DOCUMENT RESUME

ED 249 993 IR 050 861

AUTHOR	Schroeder, J. R.; And Others
TITLE	Processing and Data Distribution within the Research Libraries Information Network.
INSTITUTION	Research Libraries Group, Inc., Stanford, CA.
SPONS AGENCY	Carnegie Corp. of New York, N.Y.
PUB DATE	Nov 83
NOTE	276p.; Carnegie Study on Distributed Computing.
PUB TYPE	Reports - Research/Technical (143)
EDRS PRICE	MF01/PC12 Plus Postage.
DESCRIPTORS	Comparative Analysis; *Data Processing; Evaluation Criteria; Graphs; Library Automation; Library Cooperation; Library Expenditures; *Library Networks; Library Research; Library Services; Library Surveys; *Library Technical Processes; *Research Libraries; *Systems Development; Technology
IDENTIFIERS	*Distributed Data Processing Systems; *Research Libraries Group

ABSTRACT

Recognizing that the major technical direction of this decade is toward proximity of computing function to its users, the Research Libraries Group (RLG) is moving to distribute its processing capabilities. The goals of this study were to formulate an architecture for an RLG distributed system, to identify pre-existing components appropriate for inclusion in the network, to identify components not in existence that need to be developed, and to analyze the economic impact of distribution. The major tasks undertaken were an evaluation of candidate local systems and their vendors, using functional requirements gathered from a survey of member institutions; an assessment of current and future technology, with particular attention to availability, applicability, and price/performance over time; development of a distributed system architecture; and a financial analysis which measured the impact of distribution. The report is divided into six parts: (1) Introduction and Executive Summary; (2) A Distributed Systems Architecture; (3) The Member Survey; (4) Evaluation of Local Systems Alternatives; (5) Technology Assessment; and (6) The Financial Impact of Distribution. (THC)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

✓ This document has been reproduced as received from the person or organization originating it.
Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

David R. Richards

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

RLG DOCUMENT IDENTIFICATION
Document Code: 84-55
Brief Title: Distributed Processing
Study Final Report
Version Date: May 1984

© Copyright 1983, 1984 by
The Research Libraries Group, Inc.
Jordan Quadrangle
Stanford, California 94305

All rights reserved
Printed in the United States of America

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

J. R. Schroeder

C. F. Lerche

M. A. Ball

D. R. Richards

A. M. Tucker

J. P. Michalko

P. J. Hamblen

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

IR050861

Copyright 1983 by the Research Libraries Group, Inc.

All Rights Reserved

Acknowledgements

This work was supported by a grant to the Research Libraries Group from the Carnegie Corporation of New York. The authors wish to thank Mr. Richard Greene of Carnegie for his invaluable assistance and counsel. Thanks are also due to Karen Bendorf, Jessie Herr, Wayne Davison, and numerous others on the RLG staff for their helping hands.

Contents

Part I	Introduction and Executive Summary
Part II	A Distributed Systems Architecture
Part III	The Member Survey
Part IV	Evaluation of Local System Alternatives
Part V	Technology Assessment
Part VI	The Financial Impact of Distribution

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART I

INTRODUCTION AND EXECUTIVE SUMMARY

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Copyright 1983 by the Research Libraries Group, Inc.

All Rights Reserved

Table of Contents

Chapter 1 Introduction 1

1.1 Motivation and Goals 1
1.2 Conduct of the Study 2
1.3 Organization of this Report 2

Chapter 2 Executive Summary 5

2.1 A Distributed System Architecture 5
2.2 Trends within Member Institutions 6
2.3 Local System Alternatives 7
2.4 Technology Trends 7
2.5 Financial Impact on the Data Resource Node 8
2.6 Financial Impact on the Local Institution 9

Chapter 1

Introduction

> The center of the network is not the highly shared mainframe computer, but the user who reaches out through his network to any number of resources and information that might be necessary. Users must stop thinking of themselves as attachments to a mainframe computer.
>
> Robert Metcalfe,
> Chairman of the Board,
> 3-Com Corporation

1.1 Motivation and Goals

This study began with a conclusion: RLG must move and is moving to distribute its processing capabilities. The validity of this conclusion is demonstrated by the number of member libraries where local computer systems are planned or are already installed. RLG has already recognized that the major technical direction of this decade is toward proximity of computing function to its users. Given this assumption, the goals of the study were

- To formulate an architecture for an RLG distributed system.
- To identify pre-existing components appropriate for inclusion in the network.
- To identify components not in existence that need to be developed.
- To analyze the economic impact of distribution.

1.2 Conduct of the Study

Figure 1 portrays schematically the groups of tasks undertaken in this study, and the major interrelationships between them. The major tasks were:

- an evaluation of candidate local systems and their vendors using functional requirements gathered from a survey of member institutions;
- an assessment of current and future technology, with particular attention to availability, applicability, and price/performance over time;
- a distributed system architecture;
- and a financial analysis which measured the impact of distribution.

1.3 Organization of this Report

The report is divided into relatively self-contained parts, and is in a sequence that seems logical to the authors. At the same time, the reader may choose to use a different sequence. To facilitate "skipping", antecedent information has been included where appropriate at the expense of some redundancy. The report is composed of

Part I — **Introduction and Executive Summary** (the document you are reading), which explains the motivations and goals of the project, the organization of the report, and describes major findings and conclusions.

Part II — **A Distributed Systems Architecture**, which describes the distributed network in terms of the roles of each location, the placement of data and processing capability among locations, and the rules for inter-location communication.

Part III — **The Member Survey**, a report of findings from a survey of RLG member institutions which solicits system requirements and long-range planning information.

Part IV — **Evaluation of Local Systems Alternatives**, a description of the local system alternatives, with three candidates singled out for qualified recommendation.

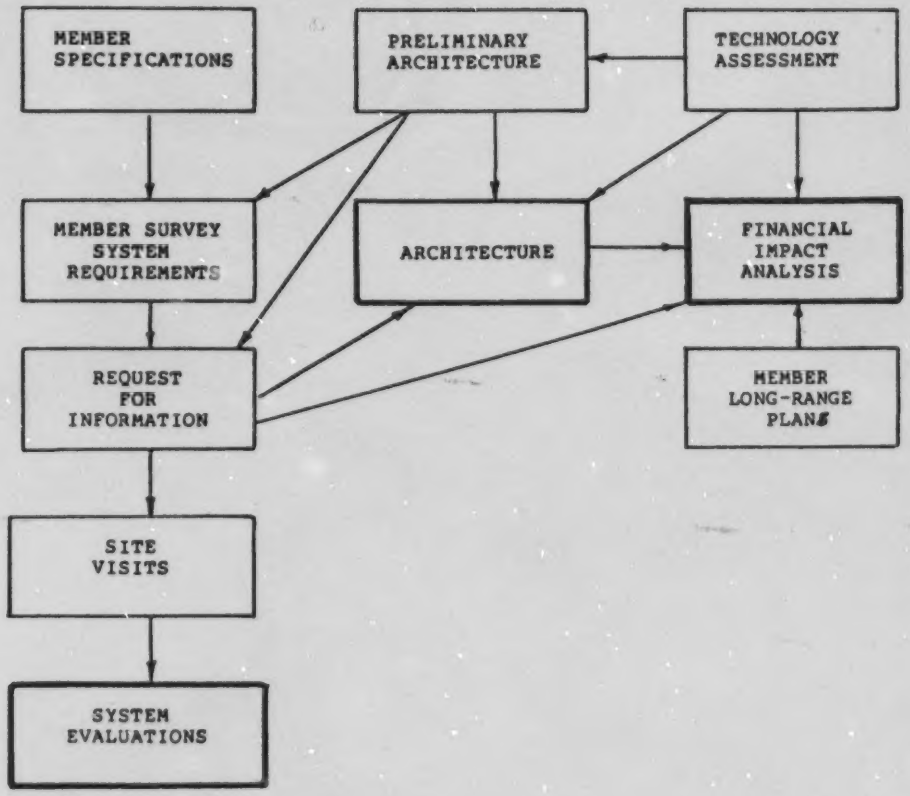
MEMBER SPECIFICATIONS
PRELIMINARY ARCHITECTURE
TECHNOLOGY ASSESSMENT
MEMBER SURVEY SYSTEM REQUIREMENTS
ARCHITECTURE
FINANCIAL IMPACT ANALYSIS
REQUEST FOR INFORMATION
MEMBER LONG-RANGE PLANS
SITE VISITS
SYSTEM EVALUATIONS

Figure 1.

Part V — **Technology Assessment**, which reviews areas of technology critical to RLG, and attempts to predict the availability and cost of new kinds of capabilities.

Part VI — **The Financial Impact of Distribution**, in which a financial model is described along with three initial scenarios and their outcomes.

Chapter 2

Executive Summary

2.1 A Distributed System Architecture

Definition of Architecture: By "Architecture" we mean a statement which defines where application processing capability is placed in a network, where data are placed in the network, what logical interconnections must exist among nodes of the network, what data are transmitted between nodes, and what kinds of inter-nodal communication protocols must be used.

Architecture Overview: The RLG distributed system is conceived of as a fully-connected network consisting of a data resource node (referred to as the central site) located at Stanford, California, and local library nodes installed in RLG member institutions.

Role of the Local Nodes: The primary role of the local node is to provide access to the library's collection through an on-line public access catalog. It will ultimately support circulation, serials control, acquisitions, cataloging, and inter-library loan. When a scholar or a library staff member is unable to locate material in the local catalog, the system will electronically transfer the search request to the data resource node. Thus, the library system serves as a "gateway" between the local area network and the RLIN network.

Role of the Data Resource Node: The primary role of the data resource node is to provide bibliographic, location, procurement, and preservation planning information covering the combined RLG collections to searchers (both staff and patron) connected to local nodes. The data resource node will continue to provide the basis for resource sharing, shared collection management and development, and cooperative preservation. In addition to its emerging role as the data resource node, the central facility will continue to offer its present services for the foreseeable future.

Local System Architecture: Although most candidate local systems reside on a single computer, the architecture is admissive of alternatives that do not. We recommend, however, that extreme care be taken with alternatives consisting of two or more heterogeneous systems lest the complexities of achieving integration outweigh other perceived advantages.

Inter-Modal Communication: Local systems must adhere to the International Standards Organization Standard 7498 - Open Systems Interconnection - Basic Reference Model and related protocol standards for all communication among RLG distributed system components. Unless this common language is employed by all local systems and by the data resource node, the distributed system, and hence the RLG programs, will not function effectively.

2.2 Trends within Member Institutions

These trends are reported from the results of the Member Survey:

Consideration of Local System Alternatives: Most RLG member libraries, even those who now operate automated circulation systems, are in some stage of local system evaluation or installation. Among this group, over twenty candidate systems are being examined. A primary goal in most of these efforts is the installation of an on-line public access catalog, but in a significant number of cases systems are being considered which are capable (or soon will be) of doing all processing, including acquisitions and cataloging.

Migration of Processing to Local Systems: According to the survey, the installation of local systems with technical processing capability will result in a 40% decline in central cataloging during the initial four years of the period, and a decline of 28% in acquisitions during the same interval. Although no further decline in central technical processing activity was indicated during the latter four years, project staff feel this to be due to local planning horizons that do not go beyond 1987. The income loss will be partially offset by search requests sent to the data resource node for bibliographic records to support these applications.

Retrospective Conversion Plans: A concomitant of on-line catalogs is a large corpus of machine-readable bibliographic records to be used therein. Members have indicated their intention to input over 6 million records over the next eight years, using RLIN as the primary vehicle. RLG will take the necessary steps to provide software and data to support this effort.

Local Area Networks: At least one-half of the RLG member institutions are actively planning for or have already installed local area networks (LANs). The resources on these networks will in all cases include the local library system, with its on-line catalog. By 1985, the average RLG LAN will have 800 devices (terminals and workstations) with access to the catalog. By 1987, the average is estimated at 1500.

2.3 Local System Alternatives

The Member Survey and the Request for Information: A summary-level requirements outline consisting of 12 sections, each covering a different functional area, was cast in survey form and sent to the RLG member libraries. The member staffs indicated their judgement of the relative importance of each functional item. The results were used to create a Request for Information, which was sent to the twenty-odd vendors whose systems were under consideration. The responses were evaluated using results from the member survey. This information was supplemented with site visits to the most promising vendors.

Local System Recommendations: There were no systems that received our recommendation without reservation, since none of the candidates support the prescribed inter-system communication protocols, and all candidates are missing minimum-level support for at least one application. Recommended with reservations are the BLIS system from Biblio-Techniques, Inc.; the GEAC system from GEAC, Ltd.; and the NOTIS system from Northwestern University. To remove the reservations, the vendor must implement the standard inter-system linkage as specified in Part II; and must complete the development of missing applications.

2.4 Technology Trends

Processor Price/performance: Processors within a given performance category have declined in price by 20% per annum over the last five years. We expect that trend to continue with respect to large computers (e.g. RLG's IBM 3081). It may accelerate beyond 20% in the case of smaller machines. This means that a constant expense for processor lease or amortization can produce a doubling of capacity every four or five years.

Magnetic Mass Storage: Currently, both local and central systems depend upon high-speed high-capacity magnetic disks to store the on-line data base. We expect the price/capacity ratio for these devices to improve at a rate of 10-12% per year for the next five years. Storage densities during this period will be improved by the introduction of new media coatings; accessing capacity will be maintained by the addition of more access mechanisms. At the end of the period, there is liklihood of a discontinuous improvement in price/performance due to the introduction of a radically different recording technique, perhaps by as much as a factor of two.

Optical Mass Storage: Optical digital storage devices have enormous potential value for RLG because of the probability of an order of magnitude advantage in storage density over magnetic devices. The reproducability, removability, robustness, and archival qualities of the media may allow the physical replication of the entire RLG data base for weekly or semi-weekly distribution and subsequent local use by the end of the decade. In addition, they may form the primary means for electronic document storage systems, with capabilities for demand publication in local or remote locations.

Long-Distance Communication: The choice is currently unclear between satellite and terrestrial circuits as a communication medium for RLG in the long term. Deregulation, the advent of fiber-optics in long-distance applications, and the conversion of voice transmission to digital signaling all suggest the latter. This is reinforced by the high capital cost, long amortization period, and operational complexities associated with a dedicated satellite network. RLG must take a strong initiative in analyzing this issue soon.

Software Engineering A great deal of attention has been focused in the last decade on "software Engineering," the processes by which computer software is designed and implemented, and the resulting innovations in methodology have had a positive impact upon software cost and overall quality.

Operating Systems: A bleak assesment is given to the current availability of operating systems appropriate for transaction and data base management. This is in marked contrast to the substantial improvements in hardware price/preformance cited above. Unless system users demand a change, and the hardware vendors recognize the existence of a potential market, this situation will hamper the development of information systems over the next decade.

2.5 Financial Impact on the Data Resource Node

The Financial Model: Project staff constructed models to predict the size of the data base and to predict computing requirements given activity estimates from the member survey. These predictions were combined with hardware price/performance projections from the technology assessment, debt service requirements, staffing assumptions, and other cost items to produce a model of RLG central site operating expense over the next eight years. Likewise, an income model was derived from the activity levels and from pricing assumptions.

Initial Scenarios and Outcomes: An inertial scenario was run assuming minimum change over and above the decline in central site activity implied by the survey. A large cumulative deficit was shown. A second scenario was run assuming conservative upgrades in processing capacity in the central site, and allowing the electronic referral of staff and patron searches from local systems to use the extra capacity. The outcome was a substantially smaller deficit. A third and more aggressive scenario assumed a large addition of capacity, and consequently more referred searching to use the capacity. The

deficit was almost totally eliminated. Subsequently, a forth scenario was run which assumed more realistic timelines for distribution together with modest price increases. This scenario produced a substantial surplus over the period examined.

Re-Orientation of the Data Resource Node: If the preceding analysis is correct, then it is clear that a new product and service profile must be sought for the data resource node if it is to survive financially. Project staff chose referred searching because of its obvious program value. There are other alternatives, however, and they need to be articulated and studied.

2.6 Financial Impact on the Local Institution

Current Operational Funding: At the gross level of analysis, current operational funding levels for automation in member libraries seem barely adequate to support initial installation and operation of local integrated systems. The picture may be better or worse in specific instances.

Local Capacity and Demand: Most members will initially support the on-line catalog using a modest number of terminals: between fifty and two hundred. The demand for more access (and hence more processing capability) will build as local area networks are installed and as patron behavior changes. Larger data storage and more processing capacity will also be necessary as a result of retropective conversion activities. In short, system capacity may lag demand for the foreseeable future. Also, a capacity deficit may exist at the outset given lack of vendor experience in sizing on-line catalog systems.

Financing Capacity Growth: The need to continuously upgrade the capacity of a local system, and the need to periodically replace worn-out or obsolete components suggests a "pay-as-you-go" philosophy for all equipment. Leasing or purchase and amortization over a four-year period is recommended. If grants are obtained for equipment purchases, then expenditures that would otherwise go for leasing or amortization should be placed in an equipment replacement pool. This approach will, of course, require the existence of capital funding. It may also require a greater institutional commitment at the outset to sustain a larger operational budget.

Institutional Costs Over Time: Assuming an "average" RLG member institution and amortization of equipment, total institutional costs for library automation are shown to increase by a factor of 2.5 over eight years. These increased costs are not all necessarily borne by the library budget, since end-user use of the DRN implies the need for allocation of funds to end-users for this purpose, and for an appropriate charge-back mechanism. Libraries should undertake now to initiate discussions within their institutions to prepare for this eventuality.

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART II

A DISTRIBUTED SYSTEM ARCHITECTURE

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Copyright 1983 by the Research Libraries Group, Inc.

All Rights Reserved

Table of Contents

Chapter 1 Introduction 1

Chapter 2 General Assumptions 3

2.1 Continuing Support for RLG's Programmatic Goals 3
2.2 Presentation of an Integrated Interface 3
2.3 Evolutionary Distribution 4
2.4 A Plurality of Local Systems 4
2.5 The Role of the RLG Central Staff 4

Chapter 3 Distributed System Overview 5

3.1 The Local Library Nodes 5
3.2 The RLG Data Resource Node 5
3.3 Regional Nodes 6

Chapter 4 Distribution and Placement of Function 7

4.1 Current Centrally-Supported Functions 7
4.2 Current Locally Supported Functions 8
4.3 Augmentation of Local Function Over Time 8
4.4 Residual Processing Support from the Data Resource Node 9

Chapter 5 Distribution and Placement of Data 11

5.1 Current Structure of the RLG Data Base 11
5.2 Data Distribution via Subset Migration 13
5.3 The Post-Distribution Data Resource Node 14
5.4 The Post-Distribution Local Library Node 16
5.5 Internodal Data Flow 19

Chapter 6 Three Architectural Models for the Local Node 21

6.1 A Heterogeneous Set of Discrete Systems 23
6.2 Single System with Multiple Applications 23
6.3 Integrated Set of Servers on a LAN 23

Chapter 7 Protocols for Inter- and Intra-Nodal Communication 29

7.1 Level 1 - Physical Link 29
7.2 Level 2 - Data Link 29
7.3 Level 3 - Network 31
7.4 Level 4 - Transport 31
7.5 Level 5 - Session 31
7.6 Level 6 - Presentation 31
7.7 Level 7 - Application 32

Chapter 8 Discussion of Intra-Nodal Data Flow 35

Chapter 1

Introduction

This document describes the architecture for a distributed RLIN system. By "architecture" we mean a statement which contains the following information:

- Where application processing capability is placed in the network.
- Where data are placed in the network.
- What logical interconnections must exist among nodes, and what their purpose is.
- What kinds of interconnection protocols will be used.

The next chapter puts forward the general assumptions that underlie the proposed architecture. Chapter 3 then presents an overview of a distributed RLIN system. Chapters 4-7, in turn, address each of the above questions in detail. Finally, the last chapter points out some of the difficulties inherent in implementing a local system as a heterogeneous collection of discrete application systems.

Chapter 2

General Assumptions

2.1 Continuing Support for RLG's Programmatic Goals

Evolution of RLIN to a distributed processing architecture must not compromise RLG programmatic goals in any way. Sufficient information will be retained and made accessible to assure the proper functioning of the Collection Management and Development Program, the Preservation Program, and the Shared Resources Program. The Shared Resources Program, in particular, will be significantly enhanced by placing a powerful existence- and location-finding tool directly in the hands of the scholar; this is not technically feasible in the current centralized system.

2.2 Presentation of an Integrated Interface

Initially, local systems may not support all applications described in this document nor all functions within them that are ultimately required. But the local system must allow data maintained and capabilities provided in one subsystem to be accessed in another. For example, circulation or serials control information should be accessible in catalog searching, and catalog searching should be available in acquisitions. The boundaries between different parts of the data base and among various functions of different subsystems should be seamless. Furthermore, a high degree of integration is desirable between the library system *per se* and other information resources in use at the institution (e.g., departmental data bases, administrative information centers, A & I services).

2.3 Evolutionary Distribution

Movement toward an RLG-wide distributed system will occur in steps over a period of several years. There will be differing timelines in each institution due to funding availability; to the variety of library computing equipment now in use; and to differing institutional plans for local area networks.

2.4 A Plurality of Local Systems

Given the number of local systems now installed and running in RLG institutions and the multiplicity of different campus computing environments, the architecture must support a plurality of local system choices. There will be no attempt to limit choice to a single alternative.

2.5 The Role of the RLG Central Staff

RLG's central staff role may undergo substantial changes in a distributed environment. In the past, the staff has provided support for the specification, design, development, and operation of computerized library information processing services. The emphasis is expected to shift in the future to include facilitation in addition to product development, particularly in the area of standards, and in the formulation of operational models. Development activities will increasingly tend toward support of wide-area networking and data synthesis and distribution.

Chapter 3

Distributed System Overview

The RLG Distributed System is conceived of as a network consisting of a data resource node located at Stanford, California, and local library nodes installed in RLG member institutions. Each node must be able to communicate with every other.

3.1 The Local Library Nodes

The primary role of the local library node is to provide access to the local collection through an on-line public access catalog (OLPAC). It also may support circulation, serials control, acquisitions, and cataloging.

The local library node functions both in the RLG distributed network and in the local network. It plays the role of a gateway. When material cannot be found locally, search requests are routed elsewhere to be processed. If an item is found, an interlibrary loan request can be sent electronically.

3.2 The RLG Data Resource Node

The primary role of the data resource node is to provide a union data base to support the resource sharing programs of RLG. In addition to supporting shared cataloging, the data resource node supports interlibrary loan with location information; supports shared collection development and management with selection and order information; and supports cooperative preservation with records of microfilming decisions.

3.3 Regional Nodes

Currently, regional nodes appear unnecessary for the RLG distributed network. RLG will move to a distance-independent communications system within the next several years. Replicating data at yet another level and staffing the regional nodes would add significantly to total network cost.

Chapter 4

Distribution and Placement of Function

This section explains the current RLIN network. It describes the migration of function into local library nodes as the linkages become operational.

4.1 Current Centrally-Supported Functions

RLIN currently supports catalog searching on a data base containing over seven million bibliographic records and two million authority records. Cataloging is the primary activity, at a level of over 1.6 million titles per annum. Acquisitions ordering, receipt, claiming, canceling are supported, but fund accounting, invoicing, and disbursement are not. Current level of use for acquisitions is 800 thousand titles per annum.

A relatively new service, the Interlibrary Loan Subsystem, allows requests to flow back and forth between borrowing and lending institutions. The data base supports this activity by providing bibliographic descriptions for the requests and location information for request routing. Current ILL activity is 109 thousand requests per annum. As the collection management and development program progresses, and as the data base is enriched with location information resulting from retrospective conversion activities, ILL use will undergo a dramatic increase.

An authority control system that will perform real-time validation of designated headings in new catalog records will be available in 1985.

The RLG central computer is an IBM 3081 Model K with 24 million bytes of main memory and 16 channels. The operating system is MVS/XA; the timesharing system is Stanford University ORVYL. The terminal communications network is supported on four PDP 11/44 front-end computers under RSX-11M and Decnet. The system supports nine transcontinental 9.6KB backbones that each break into four 2.4KB circuits at the remote end. Each 2.4KB circuit handles between fourteen and eighteen polled terminals. The number of terminals currently on the network exceeds 600. The data base is stored on twelve IBM 3380's with over two billion bytes each.

4.2 Current Locally Supported Functions

Most locally supported library automation in RLG institutions currently consists of stand-alone circulation systems. Various on-line public access catalogs are presently being evaluated, and one institution is in the process of installing an OLPAC. A growing number of RLG members have produced specifications for an integrated local system, a trend that has been observed and reported on previously. [1] [2]

Two RLG members currently operate locally developed full-function integrated systems: Northwestern University and The Pennsylvania State University. Both institutions support all technical processing functions, circulation, and an OLPAC. Penn State has nearly completed a retrospective conversion, and Northwestern has begun a selective retrospective conversion effort to enhance access to the collection.

4.3 Augmentation of Local Function Over Time

Over the next several years, member institutions can be expected to augment current local processing capabilities by adding component subsystems or by acquiring an integrated system. It seems unlikely that components or entire systems will be acquired in the future, however, unless those systems support inter- and intra-nodal linkages as described in this document. Given those linkages, RLG encourages its members to move toward a full-function local system that supports an OLPAC, circulation control, serials control, acquisitions, cataloging with authority control, and interlibrary loan.

---

1. De Gennaro, Richard. "Library Automation & Networking: Perspectives on Three Decades", Library Journal, April 1, 1983.

2. Matthews, Joseph R. "The Automated Library System Marketplace, 1982: Change and More Change!", Library Journal, March 15, 1983.

4.4 Residual Processing Support from the Data Resource Node

RLG has a two-tiered membership structure in which smaller specialized institutions participate in subject-related programs without assuming the burden of full membership. These special members, because of their generally small size, may find the idea of a local system unattractive, both for programmatic and financial reasons.

In addition, RLG provides services to a number of other libraries through a broker arrangement with CLASS (California Library Authority for System Services). These institutions may also wish to continue using centralized RLIN services longer than larger RLG member institutions.

Some percentage of RLG member institutions intend to continue using the central RLIN services for at least the next eight years, developing locally only those functions not supported centrally.

For these reasons, current central services will continue to be provided for the forseeable future.

Chapter 5

Distribution and Placement of Data

This section describes the logical structure of the current RLG data base, defines a mechanism for its distribution, and explores the post-distribution data structures in the data resource and local library nodes.

5.1 Current Structure of the RLG Data Base

The current logical structure of the RLG data base is shown in Figure 1. The circles represent entities; the arcs symbolize binary relationships among the entities. The picture does not represent physical file structure. The entities and some major relationships are described below.

- **Local Bibliographic Description.** Each cataloging user of RLIN maintains an individual bibliographic description of an edition of a given work held by the institution. The descriptions conform to U.S. MARC formats for books, serials, maps, films, and music. Associated with each description are location- and copy-level holdings statements.

- **RLG Bibliographic Description.** If any local bibliographic description exists, the one having the most complete cataloging and content designation becomes the RLG master description. An institution identifier associates the master and local descriptions. This association may indicate that the institution holds the work; that it has ordered it; and/or that it has microfilmed or will microfilm the work.

- **Order.** When an item is placed on order, the order information is appended to an existing local bibliographic description. If a description does not exist, then one is derived either from the master or from another local description; or a brief description is entered, and later upgraded or replaced by a record conforming to RLG standards for completeness. Selection decisions are also recorded here.

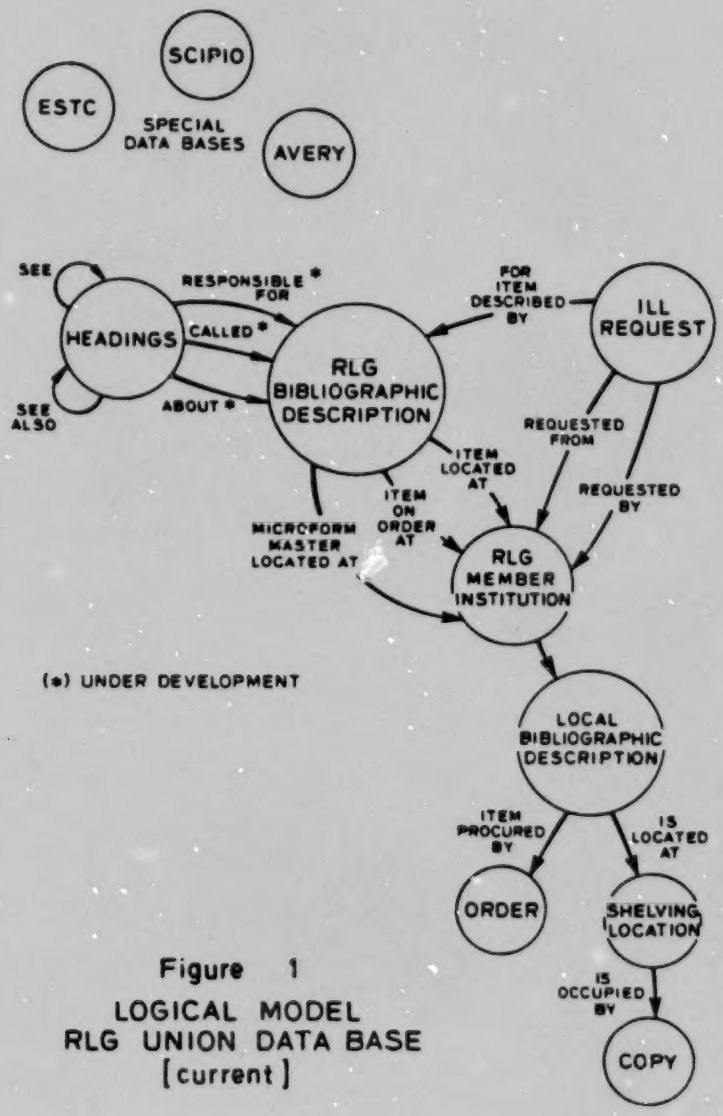

Figure 1
LOGICAL MODEL
RLG UNION DATA BASE
[current]

- **Heading.** Currently, the RLIN system supports a search-only file of headings that conform to the U.S. MARC format for authorities. RLIN will not support validation and control until sometime in 1985. The structure described here will exist after the implementation of heading validation and control. Logically, most headings will be part of a resource file, to be used for authority work and validation. In addition, any heading that has validated a record will be given a "controlling" status.

- **Interlibrary Loan Request.** Interlibrary loan requests are stored in the RLIN data base, and are transferred between institutional "mailboxes" as transactions move back and forth between borrowing and lending institutions. Associated with each request is a bibliographic description, a requesting institution, a lending institution, and local call number/location.

- **Special Data Bases.** Special data bases contain descriptions of materials for which no MARC format exists, but which have importance for some segment of the RLG constituency. Currently three such data bases are supported centrally: Eighteenth Century Short Titles, produced by the British Library, augmented with North American holdings; SCIPIO, a data base of art sales catalog descriptions of interest to museums and art libraries; and the Avery Index, which provides analytic-level information for periodicals in art and architecture.

5.2 Data Distribution via Subset Migration

Data will be distributed by migrating institutionally owned subsets into the local nodes, institution by institution, at times chosen by each member. This migration is described in more detail below.

- **Orders.** If an RLG member installs a local acquisition system that supports both procurement and fiscal functions, order information in the central data base will be sent on tape to be converted and loaded into the local data base. The central data base will, however, continue to reflect outstanding orders and selection decisions for that institution. Subsequent orders and selection decisions entered in the local system will be transmitted to the central node to allow this information to be maintained.

- **Bibliographic Descriptions.** As member institutions install OLPACs with record maintenance capability, local bibliographic descriptions, access points, and holdings information migrate from the RLG union data base to the local data base. Master bibliographic descriptions remain in the union data base, together with institution codes to indicate holding locations.

- **Interlibrary Loan Requests.** When an institution adopts a local ILL subsystem, its central "mailbox" moves to the local site. An institution's ILL requests are sent directly to its mailbox, whether it resides locally or in the RLIN central node.

- **Headings.** At the time bibliographic description records migrate to a local environment, those headings that control access points within the records are copied and moved also. The remaining copy reflects the new location where the heading resides. Subsequent changes are sent to institutions using the heading. Conversely, any heading changes made locally are reported to the central node.

5.3 The Post-Distribution Data Resource Node

Although it is certain that some RLG member institutions will continue to use the data resource node for technical processing, this section assumes complete distribution. Figure 2 provides a logical model of data structures supported in the data resource node. The entities are:

- **Bibliographic Description.** This is the master cataloging record for a given work. It has been identified from among all others residing in the network as being most complete with respect to cataloging and content designation. Each RLG bibliographic description contains identifiers to show which institutions hold the work, which have it on order, and which have microfilmed it or plan to do so.

- **Heading.** Headings carry status and location information showing whether and where they are used. Headings that are not in use are retained as resource records. Institutional identifiers also indicate which institutions have employed a particular heading to control entry point information in a local catalog record, so that heading changes can be appropriately forwarded.

- **RLG Member Institution.** Institutional identifiers indicate where a bibliographic description or heading is in use.

- **Special Data Bases.** The special data bases are unchanged as a result of distribution. They continue to be maintained and accessed centrally.

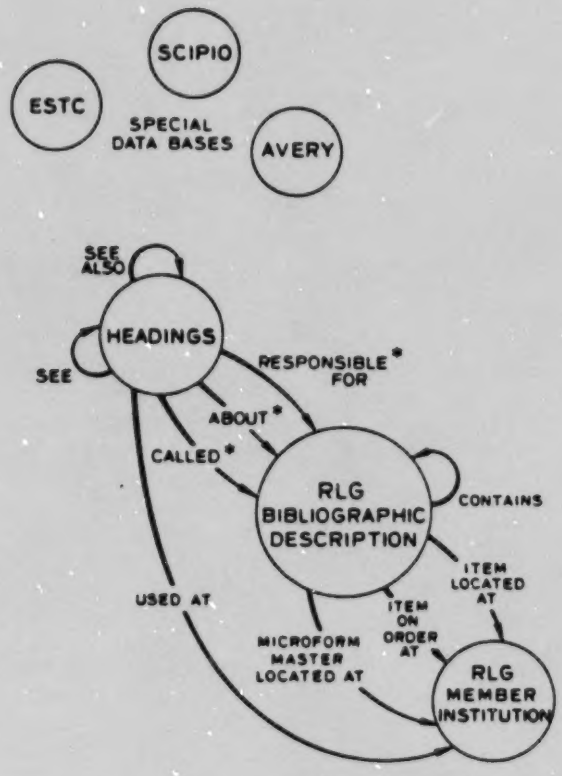

Figure 2
LOGICAL MODEL
RLG UNION DATA BASE
[post-distribution]

5.4 The Post-Distribution Local Library Node

RLG member institutions will move with varying speeds toward the installation and support of a full complement of local application subsystems. This section describes the logical data structure to be supported after all subsystems described in this document are installed at a local node. This data structure is portrayed in Figure 3. The entities are described briefly below.

- **Bibliographic Description.** All published U.S. MARC formats must be fully supported, each field in each format. In addition, supplementary local fields (local tracings, notes) and control information (identifiers and level of cataloging and content designation) are carried by the system as well.

- **Shelving Location and Physical Item.** For each title held, the system maintains a general holdings statement and call number; volume holdings, notes, and call number for each location; and notes and call number for each copy in each location. Serials control information is carried at the location level and at the item/piece level. Circulation control information is appended at the item/piece level.

- **Heading.** The published U.S. MARC authorities format is fully supported. In addition, the structure accommodates local cross references, notes, change history, heading change instructions, and appropriate control information.

- **ILL Request.** The ILL message moves back and forth between the requesting institution and potential lending institutions, picking up information as the status of the transaction changes. The message identifies the borrowing and the lending institutions, and provides a forwarding list if the material is unavailable. It contains due date/renewal information, notes, copyright and use information, and requesting patron information. Each request carries a bibliographic description of the requested item, plus local call number/location.

- **Order.** To support the acquisition functions, the system stores and maintains data for orders, requests, the full range of direct accessions, non-action selection decisions, and extended procurements. In addition, the system carries information needed to identify the supplier, to define the pieces of material in question, to properly encumber funds, to claim materials, and to route materials through in-house processing stations. Either a full or a brief bibliographic description is associated with each procurement.

- **Invoice.** The system stores and maintains supplier invoices, including supplier identification, amounts charged, currency conversion data, discounts, funds to be credited, and payment authorization.

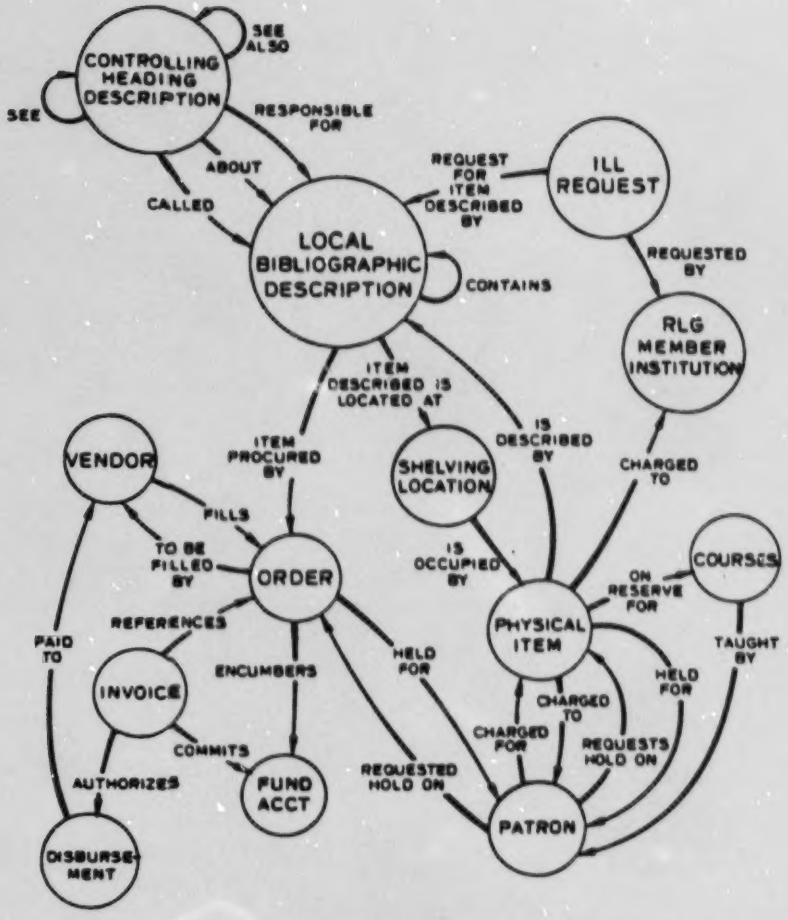

Figure 3
LOGICAL MODEL
INTEGRATED LOCAL DATA BASE

- **Fund Account.** For each fund account, the system stores and maintains account identification and descriptive information, allocation amount, encumbrances, and expenditures.

- **Vendor.** The system stores account and address information pertaining to material suppliers. Name and address information is stored for procurements, for invoices, and for claiming, and can be called out in other transactions using an identification number.

- **Disbursement.** A history of payments kept by the system includes check numbers, amounts, dates, invoice numbers incorporated in the payments, and suppliers to whom the payments were sent.

5.5 Internodal Data Flow

The previous sections indicate the need to move various kinds of data between local nodes, or between data resource and local nodes. This section explicitly defines the kind of data movement that needs to take place.

DESCRIPTION	TYPE	SOURCE	TARGET
Bibliographic Search	rt(1)	loc	DR(2)
Retrieved Record	rt	DR	loc
Bibliographic Description	ft(3)	loc	DR
Upgraded Description	ft	loc	DR
Removed Description	ft	loc	DR
Order/Selection Decision	ft	loc	DR
Decision to Microfilm	ft	loc	DR
Completion of Microfilming	ft	loc	DR
Heading Search	rt	loc	DR
Retrieved Heading	rt	DR	loc
Heading	rt	loc	DR
Modified Heading	rt	loc	DR
Removed Heading	rt	loc	DR
Modified Heading	ft	DR	loc
ILL Request	rt	loc1	loc2
ILL Response	rt	loc2	loc1

Notes:
(1): rt is a real-time transfer
(2): Data resource node
(3): ft is a batch file transfer

Chapter 6

Three Architectural Models for the Local Node

This section describes three models for the integrated local library system. These models provide an architectural framework within which particular implementations can be categorized. The architecture is admissive of all three.

Regardless of the architecture chosen for the local node, it must fulfill certain design goals.

- **Modularity and Layering.** Layering provides the capability to exchange one component for another without disturbing the surrounding system. A well-defined network interface between the local library node and the RLIN network assigns responsibility for problems and facilitates troubleshooting. In addition, it isolates the local library node from modifications to the RLIN network characteristics, and vice versa. Most of the goals described below include the concept of layering.

- **Hardware Expandability and Flexibility.** The local library node architecture can adapt to changing institutional network configurations. It allows for the addition of capacity and the replacement of discrete components with more cost-effective devices.

- **Software Maintainability and Extensibility.** Software engineering practices employed in the development of local systems should ensure that the software is well-structured, modular, well-documented, and hence, robust, maintainable, and extensible. Use of a high-level strongly typed language such as Pascal is highly desirable.

- **Standardization.** Hardware and software, especially at interconnections between the local library node, the local area network and the RLIN network, conform to ANSI and ISO standards whenever these exist.

- **Reliability.** The local library node is reliable and allows for redundant equipment for those applications demanding little or no down time. Adequate service is available to all RLG member institution locations.

- **Accessibility.** The local library node is accessible to other equipment on the institutional network.

- **Independence.** The architecture allows for independence between the local library and data resource node physical file structures.

- **Integration.** The local library node architecture allows for the possible integration of currently installed library computing equipment.

A common element of the models is the RLIN network server. The network server provides a standard interface between local node components and the RLIN network. In addition, it serves as a gateway between RLIN and the institutional network. It can also act as a local concentrator for terminals continuing to be directly connected to the data resource node. Depending on the local configuration, it may act as a router among the communicating components of the local node.

6.1 A Heterogeneous Set of Discrete Systems

This architecture provides for the interconnection of several discrete systems, each supporting a different application, together with the server described above. The data bases to support each application reside on different machines. Inter-data base searching and record transfer are required to achieve the desired levels of integration. Some terminals are local to each application processor.

Figures 4 and 5 provide schematic portrayals of two variations on this theme: one that uses a LAN as the means of interconnection, and one that employs point-to-point connections to the network server to allow message transfer among the component systems. In the LAN model, it is possible for terminals attached to LAN terminal servers to access each application system, subject to security constraints. In this model workstations can also access various systems, provided that the proper interconnection protocols are obeyed. The transition from a point-to-point configuration to a LAN is straightforward.

6.2 Single System with Multiple Applications

This alternative, portrayed in Figures 6 and 7, possesses the same characteristics as the first alternative, except that all applications reside in a single system, accessing an integrated data base. The applications processor might absorb the functions of the network server.

6.3 Integrated Set of Servers on a LAN

This model employs a set of integrated servers: the network server; one or more servers for each application; one or more terminal servers; intelligent workstations; and a server to support the integrated local data base. This alternative requires a pre-existing LAN in the institution. Figure 8 shows this architecture.

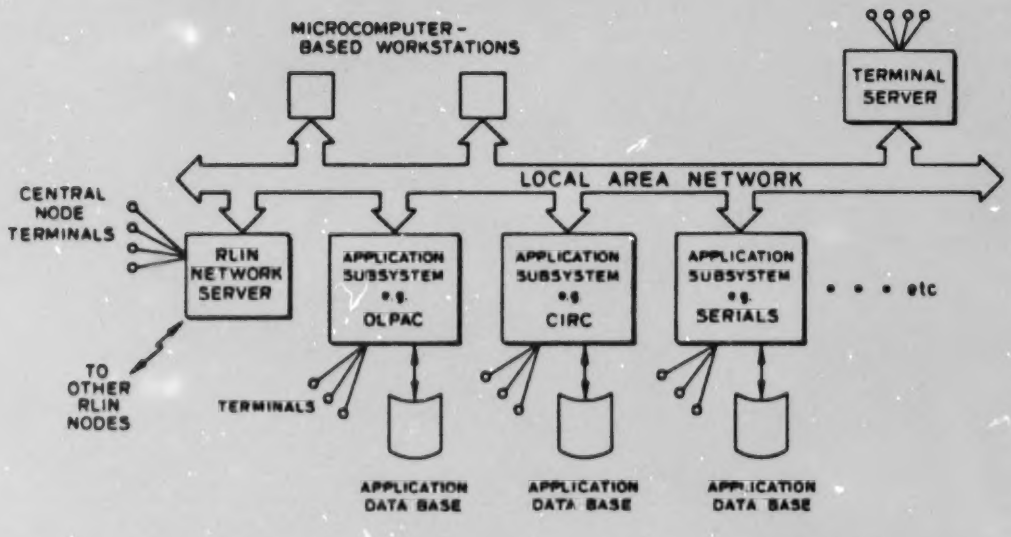

Figure 4
HETEROGENEOUS SET OF
DISCRETE SYSTEMS ON A LAN

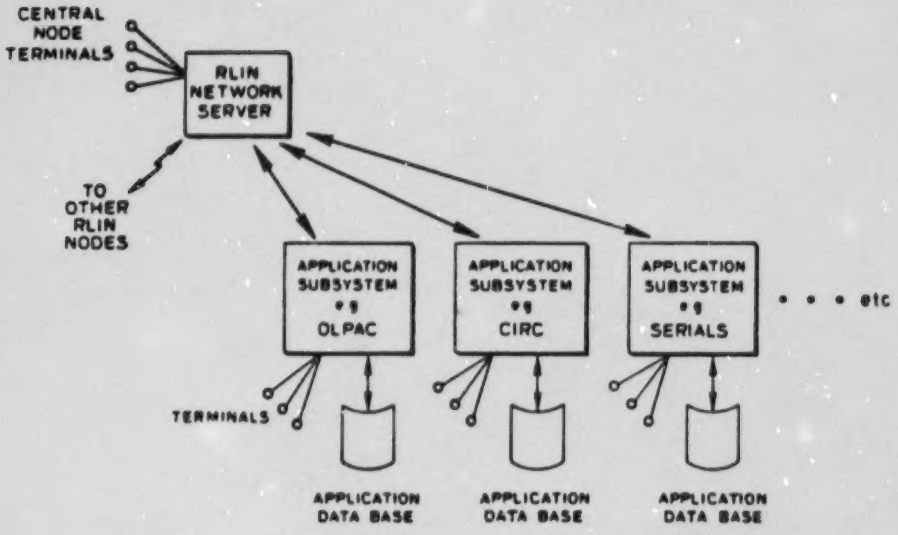

Figure 5

HETEROGENEOUS SET OF DISCRETE SUBSYSTEMS CONNECTED POINT-TO-POINT

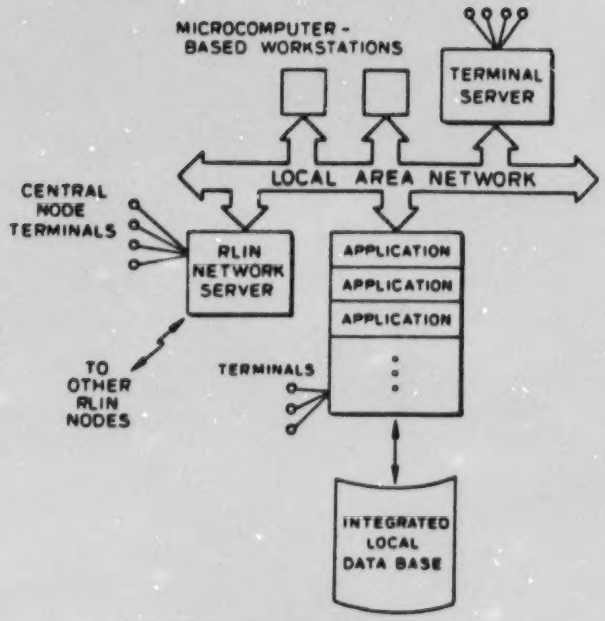

Figure 6

SINGLE SYSTEM WITH
MULTIPLE APPLICATIONS
ON A LAN

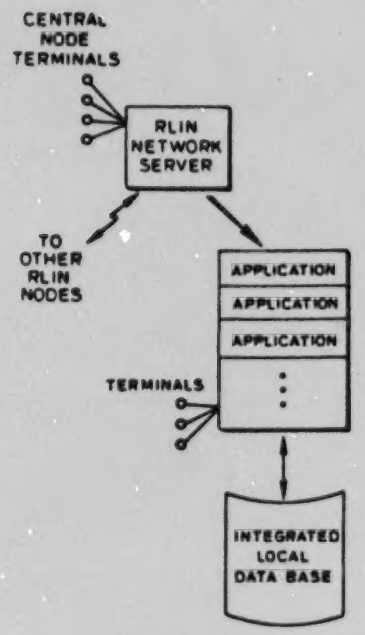

Figure 7

SINGLE SYSTEM WITH
MULTIPLE APPLICATIONS
CONNECTED POINT-TO-POINT

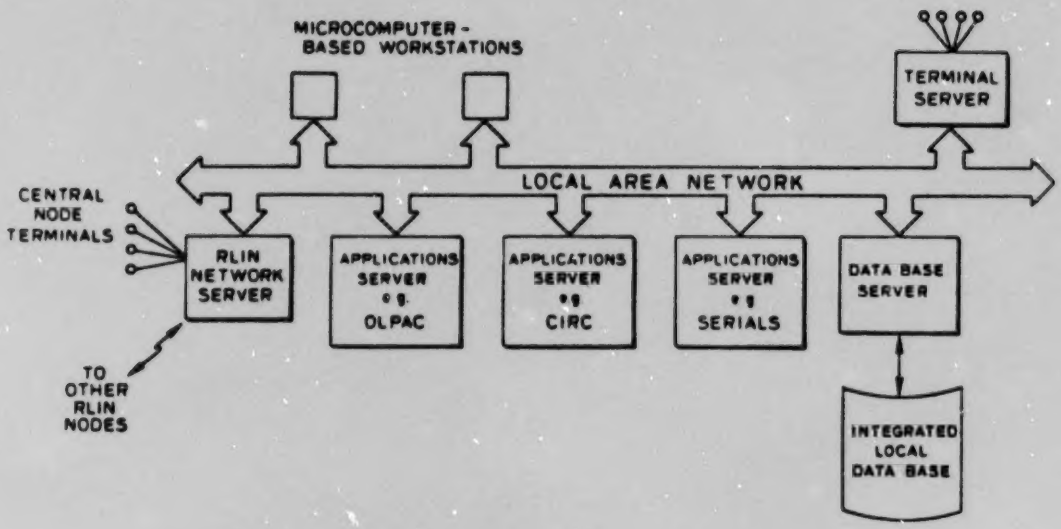

Figure 8
INTEGRATED SET OF SERVERS
ON A LAN

Chapter 7

Protocols for Inter- and Intra-Modal Communication

Local systems will adhere to International Standards Organization (ISO) Standard 7498: Open Systems Interconnection - Basic Reference Model and related protocol standards for all communication among RLG distributed system components.

The purpose of the OSI Reference Model is to identify an organized set of services and protocols necessary to achieve communication between applications independent of their location, data format, or nature of the communication. The Reference Model consists of seven hierarchically related layers, shown in Figure 9 and described in detail below.

7.1 Level 1 - Physical Link

The physical layer provides mechanical, electrical, functional, and procedural characteristics to activate, maintain, and deactivate physical connections for bit transmission between data link entities.

RLG requires support of the electrical, mechanical and functional characteristics specified in EIA Standard RS-449. Speeds will be supported up to 56 kilobits per second. Support should conform to the requirements laid out in EIA Industrial Electronics Bulletin No. 12, "Applications Notes on Interconnection Between Interface Circuits Using RS-449 and RS-232C".

7.2 Level 2 - Data Link

The data link layer provides the means to exchange data between network entities separated by a single data link.

RLG requires support for full duplex data links using HDLC Class BA protocol with options 2 and 8. This is identical to CCITT Recommendation X.25 (1980) LAPB.

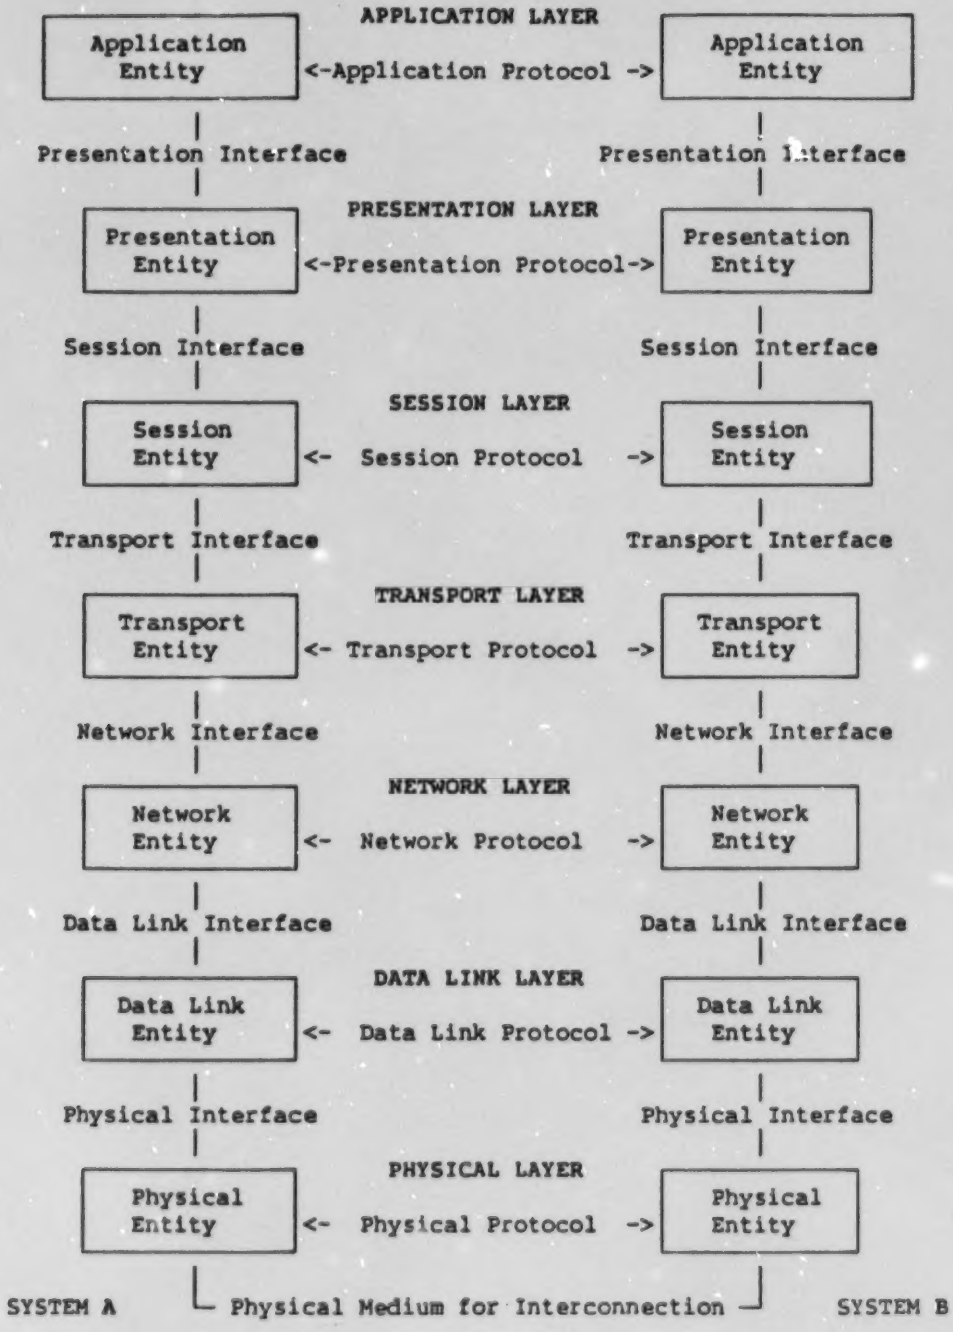

Figure 9 Open Systems Interconnection Reference Model.

7.3 Level 3 - Network

The network layer provides the means to exchange data between two transport entities. It provides transport entities with independence from routing and switching considerations.

RLG requires support for CCITT Recommendation X.25 (1980) for DTE to DCE connection to Telenet, and ISO 8208 (currently a draft proposal) for DTE to DTE connection over leased communication channels.

7.4 Level 4 - Transport

The transport layer transparently transfers data between session entities and relieves them from any concern with the detailed way in which reliable and cost-effective transfer is achieved.

RLG requires support for ISO 8073 (currently a draft proposal) conforming, at a minimum, to Class 2 using flow control, normal 7 bit format, and a TPDU length of 128 octets. Full support of Class 4 is desirable.

7.5 Level 5 - Session

The session layer provides the means necessary for cooperating presentation entities to organize and synchronize their dialogue and manage their data exchange.

RLG requires support for ISO 8327 (currently a draft proposal) conforming to the basic combined subset with "expedited data" and "use of duplex" function units in addition to the kernel. Support for segmenting and reassembly and for concatenation is not required.

7.6 Level 6 - Presentation

The presentation layer represents information to communicating application entities in a way that preserves semantics while resolving syntax differences.

In the absence of a protocol standard for this layer, RLG is currently implementing a basic presentation service. It provides for the transparent transfer of 8-bit data. RLG requires support for the ISO Presentation Protocol when it is defined; a draft is expected by the end of 1983. RLG also requires support for transfer syntaxes appropriate to the exchange of data between bibliographic systems and between such systems and those of material vendors. Such syntaxes will probably be defined as ANSI Committee Z39 standards. Extended character set support is required for the Roman, Cyrillic, Hebrew, and Arabic alphabets, and for Chinese, Japanese, and Korean characters. The bibliographic character set for each of these is substantially larger than that used for business data.

7.7 Level 7 - Application

Protocols of this layer directly serve an end-user application process by providing the information communication service appropriate to the application, to the management of the application, and to the management of the system in which it resides. Application processes are the ultimate sources and recipients of the data exchanged.

RLG is currently a participant in the development of standards for this layer. In cooperation with the Washington Library Network and the Library of Congress, and with funding from the Council on Library Resources, RLG is now implementing an application protocol for library and information science applications. This protocol provides for the establishment of a connection between bibliographic applications, for searching of remote data bases and retrieval of records, and the transfer of batches of bibliographic records in MARC format between applications. [3] [4] [5] [6].

---

3. LSP/AI-3: Application Service Definition and Protocol Specification for Record Transfer, Research Libraries Group.

4. LSP/AI-4: Application Service Definition and Protocol Specification for Information Retrieval, Research Libraries Group.

5. LSP/AI-5: Common Application Service Definition and Protocol Specification, Research Libraries Group.

6. LSP/SNI-8: Message Transfer Application Service Definition and Protocol Specification, Research Libraries Group.

Subcommittee D of ANSI Committee Z39 is currently reviewing this protocol for standardization, with extensions to support additional library functions, e.g., interlibrary loan, circulation status inquiry.

In the future, RLG will also require support for the ISO Common Application Protocol, for the ISO File Access Application Protocol, and possibly for the ISO Virtual Terminal Application Protocol as they are developed.

Chapter 8

Discussion of Intra-Modal Data Flow

The flow of data among application subsystems is straightforward in an architecture where the data base is physically integrated and under control of a single body of software (see Figures 6, 7, and 8). The balance of this discussion deals with the "Heterogeneous Systems" model, where neither of these conditions is met. Whatever physical data structure has been used to implement the logical data base model shown in Figure 3, software facilities will probably exist to support automatic maintenance of relationships among entities in real time or via deferred processing, whenever the relationships cross physical system boundaries.

When the data base is segmented across several dissimilar systems, facilities to allow one system to interrogate the data base of another are substantially more complex. Standard protocols must be supported to allow application-to-application communication across machine boundaries (see the previous section).

A further complication arises from the need to maintain relationships among entities stored in different systems. For example, one must send a procurement ID from the acquisitions system to the catalog when an order has been placed, so that an OLPAC user can know that an item is being acquired. Again, the standard protocols must be employed that support record updating across systems.

Under the "Heterogeneous Systems" model some file redundancy and batch synchronization of the two files will undoubtedly be necessary, even when boundaries are optimally placed. Of all the possible models, this one seems the least attractive.

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART III

THE MEMBER SURVEY

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Table of Contents

Chapter 1 Introduction 1

Chapter 2 Functional Requirements 3

2.1 Intended Uses 3
2.2 Methodology 3
2.3 Degree of Response Diversity 4
2.4 Survey Results: Selected Comments 6

Chapter 3 Information to Support Long-Range Planning 9

3.1 Timeline for Local Implementation 9
3.2 Homogeneity Within the Local Node 11
3.3 Estimated Activity Levels 11
3.4 Availability of Operational Funding 14
3.5 Retrospective Conversion Plans 15
3.6 Local Area Networking Plans 15

Chapter 1

Introduction

This part of the report deals with the Member Survey. The purposes of the Member Survey were twofold. The first was to arrive at a set of summary-level specifications for a local library system that would function within the RLG Distributed Computing Network. The second was to solicit long-range plans from member institutions.

The target population was RLG full member institutions, and those special members and users of the system that wished to participate. In this version of the final report, only the responses of full members are considered.

Accordingly, the remainder of this part consists of two chapters: Chapter 2 deals with the Functional Requirements portion of the survey. Chapter 3 covers the solicitation of long-range plans.

Chapter 2

Functional Requirements

Members were requested in this portion of the survey to individually rate several hundred requirement statements using a simple set of coded responses.

2.1 Intended Uses

The primary purpose of the Functional Requirements section was to allow member library staff to comment on a draft set of functional requirements, and to provide a rating of relative importance for each requirement. The information gained thereby informed the architectural design process; provided the basis for the construction of a vendor Request for Information (RFI); and allowed the RFI responses to be evaluated for a "best fit" with respect to system capabilities.

A secondary objective of the survey was the creation of a set of functional requirements which could be further improved upon and used by member institutions as part of their own system selection process.

2.2 Methodology

As a first step, project staff gathered together all extant Requests for Proposal and other relevant functional requirements definitions from member institutions. This material was used to synthesize a single requirements outline, consisting of over five hundred items divided into twelve sections. Each item elicited a coded response of the form:

- "R" meaning that the item is required, absolutely necessary.

- "D" meaning that the item is desireable, not absolutely necessary but of some use.

- "N" meaning that the item is not necessary and/or desireable.

- "?" meaning that the statement for the item is not clear, is not understood.

The survey participants took approximately six weeks to complete their responses. In general, the twelve sections were individually assigned to the different organizational units in each library to allow work to proceed simultaneously.

Project staff implemented an online data entry program for conversion of survey responses into machine-readable form. These data were then cross-tabulated by the Statistical Analysis System (SAS) to produce counts of the form #R,#D,#N,#? for each item in the functional requirements section. [1]

2.3 Degree of Response Diversity

One of the initial concerns about the survey was a possible tendency toward indiscriminate use of the "Required" designation. "R" was in fact used a great deal, because the items in the survey were drawn from member-articulated requirements to a great extent. Figure 1 shows the degree of unanimity in regard to "Required" items. It is seen that only twenty questions produced 100% "R" responses (twenty-two institutions counted); one question produced 100% "non-R" responses; and there were over seventy questions that had twenty "R"s.

From another point of view, it should be asked to what extent these responses represent a consensus, and if so, how strong. The fact that 75% of the questions evoked responses that were more that half "R"s points toward an affirmative answer, particularly when the remaining responses to those same questions were largely "D"s.

---

1. RLG Central Staff will make copies of the cross-tabulated results available to member libraries upon request.

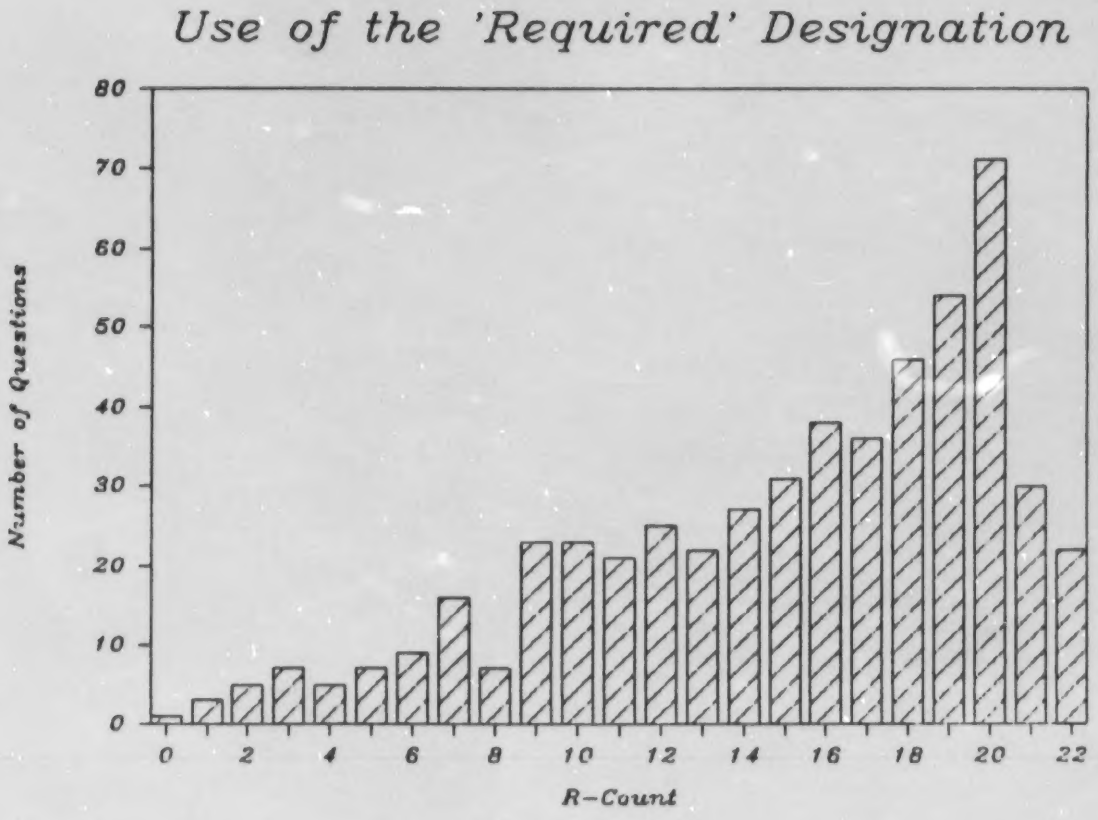
Use of the 'Required' Designation
Number of Questions
80
70
60
50
40
30
20
10
0
0
2
4
6
8
10
12
14
16
18
20
22
R-Count

Figure 1.

Finally, it is of interest to note the use of "?", indicating lack of understanding perhaps due to a poorly articulated question. One hundred twenty-seven questions had at least one "?". The average number of "?" responses per question within that subset was 1.4; the maximum was six. It was obvious in several cases that the questions were at fault, not the respondents.

2.4 Survey Results: Selected Comments

A majority of the survey results can be termed "unremarkable" in the sense that they were expected. The comments included herein are directed toward items of special importance, items which produced unexpected results, or items that carry a future impact.

The functional requirements survey validated the most basic design assumptions: that the local system should be composed of an integrated set of components, and that the basic device for user interaction should be a full-screen CRT terminal (18R). Support was not strong for the line-by-line terminal as a secondary device (12R). This is surprising, given the need to proliferate patron access terminals and to do it cheaply.

Respondents felt strongly that the primary terminal should support the ALA character set, but the support for a local system that encompasses non-roman and CJK support appeared quite weak (9R). Since it is unlikely that vendor-supported systems will provide these functions, central RLIN must continue to support them for the indefinite future.

With respect to the on-line catalog, respondents clearly want a system with both a command mode and a well-developed tutorial capability. The tutorial capability should provide both extensive prompting and context-specific help. There was moderately strong but not unanimous support for word indexing as an option (18R), and Boolean expressions (17R). It was surprising that the capability to limit results using non-indexed terms such as date and country of publication got a "required" from less than half the respondents. The current literature places a high priority on this capability. [2].

Those surveyed evidently place a great deal of importance on the design of display screens for an on-line catalog, and wish to have the ability to define new or modify existing ones (20R). Also strongly felt is the need to show circulation (19R) and acquisition (18R) status to patrons searching the catalog. This implies the need for some integration across the on-line catalog, circulation control, and acquisitions functions.

---

2. CLR study on On-Line Public Access Catalogs

There was a marked lack of enthusiasm (3R-9R) for a patron being able to request a piece of material from the terminal by entering an ILL request, placing a hold, or suggesting a purchase. Evidently, much thought and discussion must take place before this can be envisioned on a wide scale.

The results were largely unremarkable in those areas concerned with technical processing. Worthy of mention, however, is the response to a requirement for real-time heading validation: (15R) is by no means unanimous. Receiving stronger support were those editing functions needed to guarantee authority file integrity: (16R-19R).

In the System Administration arena, there were some who did not feel the need to identify and restrict users on the basis of user location (i.e. where the terminal is or is not). Given the wide geographical area over which access to an on-line catalog will be provided, we think that such an ability will prove important.

The system availability expectations of the respondents were high. Given a sixteen-hour weekday and a twelve-hour week-end day, the average expected ratio of up-time to scheduled time was 98.8 percent. The average expected MTBF was three months; average MTTR was three hours on weekdays and 6 hours on week-ends. Our experience with data centers of this kind leads us to comment that these expectations are very much on the high side of "probable".

Response expectations were also of interest. Respondents expect that 75th %-tile response for all commands should be under four seconds, and under six seconds for multi-term searches. This is, if anything, an overly modest expectation. The expectation for the 95th %-tile was seven and ten seconds for all commands and multi-term searches, respectively. To expect such consistent behavior at the 95th %-tile is not realistic; fifteen and thirty seconds would seem more reasonable, especially if the system is performing as indicated at the 75th %-tile.

Project staff took special interest in those functions dealing with local/central data movement, in whichever application or functional area they happened to fall. The reason for the interest is obvious: **without the requisite level of local/central data movement, a distributed RLIN must be counted as a failure**. There was virtually unanimous support for the ability to extend on-line catalog access beyond the local collection and into the data resource node. Inexplicably, the capability to display summary location and acquisitions information across the entire membership did not fare as well (10R). More disturbing is the lack of unanimity on the need to feed cataloging (17R) and new authority work (14R) back to the data resource node. A high priority should be given to clearing up any misunderstandings between central and library staff regarding these critical functions.

Chapter 3

Information to Support Long-Range Planning

Members were requested in this portion of the survey to provide as much information as possible regarding the library's schedule for supporting application subsystems locally; estimated system activity levels over the coming eight years; the availability of local funding; efforts to retrospectively convert the card catalog to machine-readable form; and plans to install local area networks.

It was known beforehand that well-defined plans in some of these areas were not yet in existence. Nonetheless, the objective was to gain a measure of member intent at the time of the survey. A continuing objective will be to incorporate better information into this process as it becomes available.

3.1 Timeline for Local Implementation

The survey requested planning information relative to each application that could potentially run locally. Of special interest were cataloging and acquisitions, and the planned dates by which the movement of processing to a local system would take place.

Figure 2 portrays acquisitions and cataloging titles processed centrally over the next eight years given local plans as indicated in the survey responses. It is seen that cataloging titles diminish from 1.3 million to something over 800 thousand over a four year period. Acquisitions titles fall proportionally over the same period.

Looking ahead to financial consequences, an optimist might comment that the damage is done in the first four years of the period, that it probably isn't fatal, and that slippage in plans to install local systems is inevitable.

A pessimist would point out that planning horizons do not typically extend beyond four years, and that the later four year period will see a continuation of the trend set previously. The trend, incidently, would cause a cessation of central technical processing in 1992. It is therefore of crucial importance that members keep one another and the central staff appraised of their evolving plans.

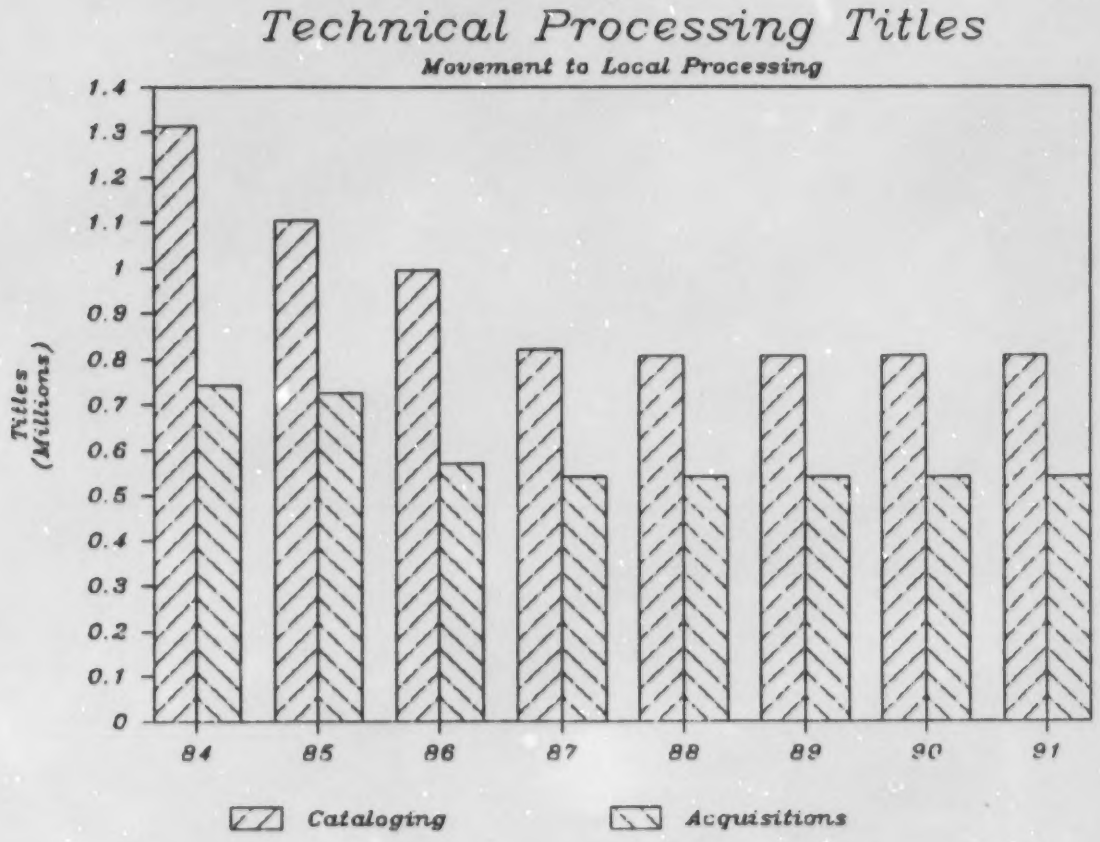
Technical Processing Titles
Movement to Local Processing
Titles (Millions)
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
84
85
86
87
88
89
90
91
Cataloging
Acquisitions

Figure 2.

3.2 Homogeneity Within the Local Node

It is worthy of comment that in some member institutions, different sets of alternative systems seem to be under consideration to support different applications. The survey provides an indication rather than proof of this. The indication lies in the degree of similarity or difference in the lists of vendors under consideration to support each application within an institution. Applying some common sense in the interpretation of "similarity" and "difference", we find nine institutions who could (at least in theory) install three or more discrete systems.

3.3 Estimated Activity Levels

The following summarizes the projected activity levels within the institutions surveyed. These were subsequently used in the Request for Information (see Part IV) to give responding vendors some quantitative profiles of the RLG member libraries. The numbers given are (small, average, large) taken across the membership.

Circulation

180K, 690K, 2000K	Annual number of charges
21K, 83K, 200K	Average number of pieces simultaneously in circulation
12K, 53K, 90K	Number of persons holding library privilages
0%, 2.1%, 7.5%	Expected per annum rate of change in circulation activity

Cataloging

13K, 45K, 110K	Number of titles cataloged annually
5%, 16%, 30%	Percentage of original input
1K, 5K, 12K	Number of added copy/added volume changes to holdings
1K, 9K, 15K	Other modifications to bibliographic descriptions

Acquisitions

7K, 30K, 107K	Number of items ordered annually
3K, 32K, 116K	Number of items received as direct accessions annually
50, 227, 700	Number of active fund accounts maintained
0%, 1.3%, 4%	Annual expected increase in acquisitions activity

Serials Control

6.5K, 35K, 140K	Current number of active serials
43K, 152K, 300K	Number of pieces checked in annually

Catalog Searching

Current searching estimates on card catalogs in RLG institutions are somewhat suspect, given the difficulty of gathering hard data. They are reasonable consistent with those few instances where direct measurement has been done.

1K, 8.7K, 20K — Number of catalog searches per day

5 — Peak/average ratio, hourly rate[3]

Estimates of the annual growth in searching are even more suspect. The following may be used to gain a rough idea of possible growth patterns.[4] It reflects institutional plans for local area networks, and the number of devices that will have access to the on-line catalog.

3%, 1054, 4000 — Devices installed by September, 1985

41, 2378, 12000 — Devices installed by September, 1987

On-line Catalog Size

134K, 487K, 1200K — Number of bibliographic descriptions now in machine-readable form

25K, 514K, 2500K — Titles that will undergo retrospective conversion over the next eight years

---

3. Taken from a Peat, Marwick, and Mitchell study conducted at New York University. This number represents a single institution, rather than the entire membership.

4. Taken from the responses of eleven RLG institutions to a member survey questionnaire on local area network planning.

3.4 Availability of Operational Funding

The following information reflects responses given to questions on sources and allocation of funding for local library automation activities. In general, the responses to this section were neither as complete nor as consistent as those to other sections. This report confines itself to those items for which complete, consistent information was obtained.

- Current local library automation budget: minimum $25,000; mean $255,000; maximum $700,000.

- Annual budget devoted to RLG for fees, assesments, and services: minimum $11,000; mean $225,700; maximum $433,655.

- Expected percentage per annum increase in local automation expenditures: minimum 2%; mean 7.8%; maximum 10%.

- Only six respondents answered completely and consistently the question which requested a breakdown of the automation budget. The remainder either left the item blank, or indicated that the information was not available. Of the six, only one institution reported any significant equipment amortization.

- The sources of funding for capital equipment were said to be as follows:

  * Library operating budget: 5 respondents

  * Centralized funding pool: 3 respondents

  * Both of the above: 9 respondents

  * Loans from a restricted endowment: 1 respondent

- Half of the answering respondents (8 out of 14) said there was no centrally-funded automation/networking activity. Of those that remained, three said they receive allocations from this budget. A larger number (11) indicated some participation in institution-wide automation planning.

3.5 Retrospective Conversion Plans

The information from the survey regarding retrospective conversion plans has been overtaken by events. Since the survey was taken, the board has agreed upon a RECON strategy which is RLIN-based, and which may change the timing and the magnitude of the activity. Secondly, the estimates provided by the survey for the years 1983-84 have not been borne out by reality. For purposes of the financial impact study, projected RECON activity for the year 1983 has been deleted, that for 1984 scaled down significantly, and the survey estimates for years 1984-91 have each been moved out one year. Figure 3 portrays the survey results. Central staff plans to update this with a new solicitation of planning information during the first quarter of calendar 1984.

3.6 Local Area Networking Plans

Twelve institutions responded with firm indications that they are either planning for the installation of a local area network (LAN), or are in the process of installing one at this time. LANs are expected to have a profound effect upon local library systems, because of the wide access they provide to the on-line catalog. Indeed, the planned existence of an on-line catalog provides much of the impetus behind some of the networking efforts.

The following table provides an idea of the projected growth over four years in the number of terminals and workstations having access to the on-line catalog in those institutions who responded to this section:

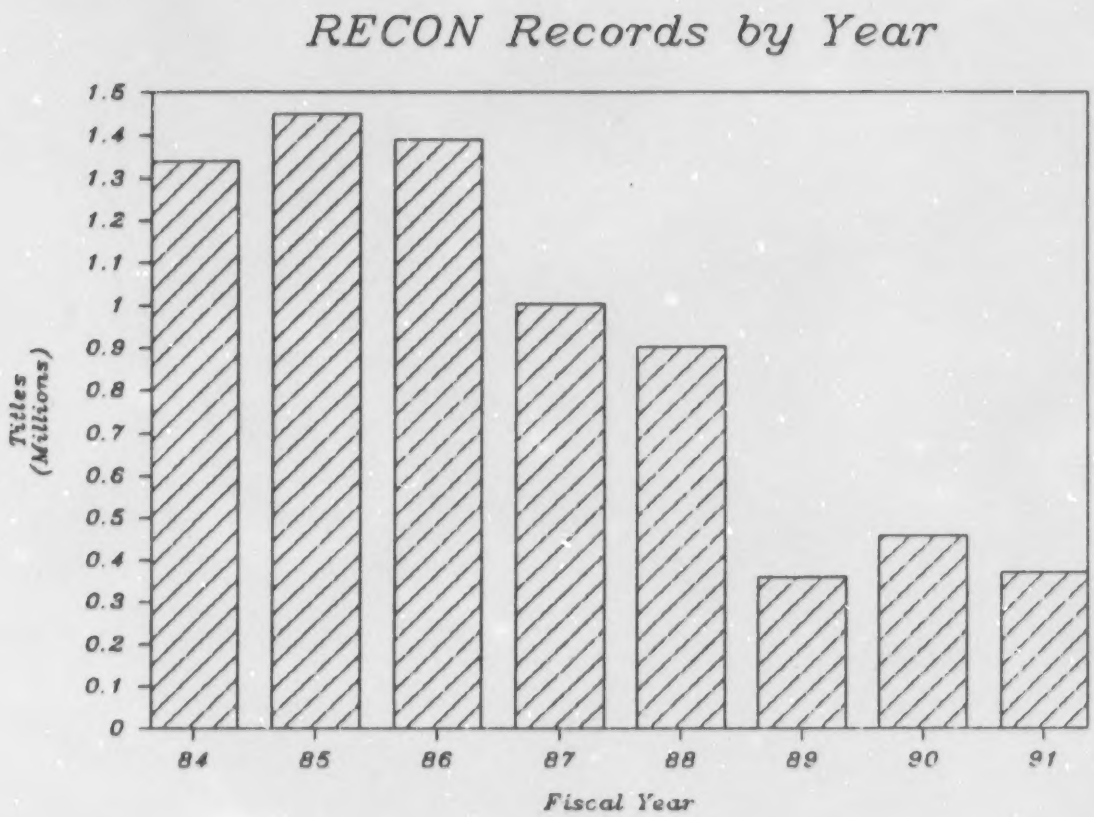
RECON Records by Year
Titles (Millions)
1.5
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
84
85
86
87
88
89
90
91
Fiscal Year

Figure 3.

Institution	Devices, 1985	Devices, 1987
1.	2250	4125
2.	250	440
3.	100	750
4.	150	625
5.	1500[*]	2625[*]
6.	1875[*]	4500[*]
7.	400	650
8.	190	190
9.	400	1100
10.	1500	1500
11.	4000	12000
12.	350	500

[*] Project staff computed these numbers by making the assumption that 75% of devices on the network will have access to the library computer.

These numbers have a disturbing quality, given the the lack of applicable experience with some of the candidate systems, the unknown behavior of current users of the card catalog in most instances, and the liklihood that the poorly understood behavior of the current user will change in major ways given the power and accessability of an on-line catalog connected to a LAN.

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART IV

EVALUATION OF LOCAL SYSTEM ALTERNATIVES

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Acknowledgements

The RLG staff wish to thank the organizations who responded to the Request for Information, and who extended their hospitality to us and took their valuable time to assure the success of our site visits.

Table of Contents

Chapter 1 The Request for Information	1
1.1 Design of the RFI	1
1.2 List of Recipients and Respondents	2
1.3 Site Visits	2
1.4 Maintaining the Currency of this Report	3
Chapter 2 Evaluation Summary	5
2.1 Recommended without Reservation	5
2.2 Recommended with Reservations	6
2.3 Partial Solutions	7
2.4 Judgement Withheld	7
2.5 Rejected Alternatives	8
Chapter 3 Summary Descriptions of the Systems	9
3.1 Avatar	10
3.2 Biblio-Techniques	14
3.3 Carlyle	21
3.4 CLSI	26
3.5 GEAC	31
3.6 Innovative Interfaces	40
3.7 NOTIS	42
3.8 RLIN	50
Appendix A RLG Member Activity Profile	55

Chapter 1

The Request for Information

Given a preliminary statement of architecture (see Part II), and the results of the Member Survey (see Part III), project staff formulated a Request for Information (RFI) to be sent to those organizations having candidate local systems or local system components. The purpose of this instrument was to establish a body of basic information about each system in a common format to facilitate comparison.

1.1 Design of the RFI

In designing the RFI, as was the case with the Member Survey, staff gave a high priority to minimizing the time required to respond by using questions having one word answers, and allowing the use of existing materials where possible. The contents of the RFI were as follows:

- Architecture and Linkage: A draft of the proposed architecture, together with instructions to provide a statement describing the product in terms of one of the hypothetical models provided. Vendors were also instructed to comment on the proposed standard for linkage protocols in terms of their potential willingness to implement support for them.

- Hardware and Software: Request for a component list of hardware and software components used in the system; prices for each component; and sizing formulae used to configure a system to handle a given load.

- Installation, Reliability, and Support: Solicitation of information relative to installation, training, documentation, site preparation, and environmental requirements; on reliability experienced in the field; and on policies and practices to provide ongoing support for the system after installation.

- System Operation and Administration: Solicitation of information concerning operational and administrative procedures and requisite staffing levels.

- Satisfaction of Functional Requirements: Solicitation of information relating to conformence to functional requirements determined from the member survey.

- Vendor Qualification: Request for an installed system list; organization and staffing level information; and a statment of earnings and current balance sheet.

- Costed Configurations: Vendors were asked to submit costed configurations for three hypothetical libraries: small, average, and large. Definitions for these were formulated from the Member Survey responses. See Appendix A - RLG Member Activity Profile.

1.2 List of Recipients and Respondents

The following lists those organizations who received a copy of the RFI. They are categorized according to who responded and who did not.

RESPONDENTS	NON-RESPONDENTS
Northwestern	Faxon
GEAC	University of Guelf
VPI/Hewlett-Packard	CTI
CLSI	Penn State
BiblioTechniques/WLN	DataPhase
University of Wisconsin	Blackwell NA
Carlyle Systems	BRS
Innovative Interfaces	Ringgold
Avatar	
RLIN	
IBM	

1.3 Site Visits

Project Staff conducted site visits to all responding vendors with the exception of VPI/Hewlett-Packard, Wisconsin, and IBM. In the case of the first, their submission had no material on linkage or architecture, nor did it have a costed configuration. We will follow up in the near future to attempt to fill in the gaps. The latter two vendors were rejected as alternatives because of their specification of IBM/SNA as the linkage vehicle without first commenting in detail on our proposal and how they would implement it. We considered this unresponsive.

The site visits themselves were not only useful; they were necessary. It is difficult to draw functional specifications that can measure ease of use; we had to be able to try the products in order to form opinions about efficacy of the user interface. Secondly, it was instructive to meet the management of these organizations and to talk to them and their chief designers. Subjective impressions that are formed in such discussions should not be ignored.

1.4 Maintaining the Currency of this Report

It is the intention of the central staff to provide continuing monitoring of the systems covered in this document, and to reissue the document periodically when the degree of change warrants. This means that systems can be expected to move from one category to another as time passes.

Chapter 2

Evaluation Summary

N.B. It is emphasized that previous evaluation reports were interim in nature. Those who had access to the preliminary reports are advised to re-read this report in full.

Several categories are defined below together with the systems that fall into each category and an explanation of the placement when that is deemed appropriate. Systems are listed in alphabetical sequence within each category, with no ranking implied. Each system has strengths and weakness which should be traded off in light of local circumstances.

2.1 Recommended without Reservation

The criteria for inclusion in this section are as follows:

1. The system supports all local applications in a fully integrated fashion, and
2. the system is running in production in at least one location, and
3. the system supports the inter-system linkage protocols defined in Part II, and
4. the responding organization is willing to work with RLG over time to redress any perceived limitations.

At the present time, there are no systems in this category.

2.2 Recommended with Reservations

The criteria for inclusion in this category are as follows:

1. The system supports a critical mass of integrated local applications, and
2. the system is now in production in at lease one location, and
3. other local applications not now supported are under development, with projected completion dates no more than eighteen months in the future, and
4. the responding organization is willing to commit to the implementation of an inter-system link as defined in the RFI, and
5. the responding organization has expressed willingness to work with RLG over time to redress any other perceived limitations.

The systems in this category are:

- BiblioTechniques
- GEAC
- NOTIS

2.3 Partial Solutions

The criteria for inclusion in this category are as follows:

1. The system supports at least one local application, and

2. the system is in production in at least one location, and

3. the responding organization is willing to commit to the implementation of an inter-system link with the RLIN data resource node, and

4. the responding organization is willing to commit to the implementation of any intra-system links necessary to achieve proper integration with other local system components.

There is presently one system in this category:

- Innovative Interfaces

2.4 Judgement Withheld

These systems exhibit some promise as future alternatives, but are placed in this category because one or more contingencies must be removed before a recommendation can be made. An example is the case in which substantial development must be done before a demonstrable system exists. The systems in this category are

- Avatar. Avatar is in the midst of negotiations with OCLC that could result in Avatar being acquired, and the system becoming LS-2000.

- Carlyle. At the time of the site visit, only a demonstration system for an on-line catalog could be shown. We will revisit Carlyle after the first system is in production.

- CLSI. CLSI does not currently support full MARC records. Support for full MARC, plus upgrades to other parts of the system will occur in mid-1984. We will re-examine the situation at that time.

- RLIN. It is unclear without further feasibility studies what interest may exist on the part of a subset of the RLG membership to approve and fund the development of local applications.

2.5 Rejected Alternatives

The systems in this category are not considered to be viable alternatives at this time. They are:

- IBM. The response did not address use of the Open System Interconnection Reference Model. It proposed the use of IBM/SNA, with an X.25 interface which allows two SNA stations to communicate over an X.25 network.

- University of Wisconsin. The response did not adequately address the use of the OSI Reference Model.

Chapter 3

Summary Descriptions of the Systems

3.1 Avatar

Avatar Systems, Inc., based in Potomac, MD, is marketing an integrated library system that is an outgrowth of the ILS system developed at the National Library of Medicine.

Architecture

The Avatar system is a single system with multiple applications.

Linkage Considerations

Owing to the negotiations now under way between Avatar and OCLC, discussion of possible linkages have not been undertaken. We will proceed with these discussions should the situation warrant.

Underlying Software

The Avatar system runs under MIIS, a version of the MUMPS operating system marketed by Meditech, of Cambridge, MA. (Although there is an ANSI standard for MUMPS, MIIS does not conform to it.) This is a standalone operating system supporting an interpreted text manipulation language with an embedded hierarchical data base language; there is no mechanism for interface with code written in any other programming language.

The facts that the MIIS language is completely interpreted and that the maximum size of a single program is 2K bytes (applications are composed of multiple chained programs) encourage a coding style that uses very short variable identifiers, relies on side effects, and is inherently unstructured. This, coupled with the fact that there are very few programmers who know MUMPS creates long-term maintenance problems for Avatar. They are considering the reimplementation of the system in C, using UNIX as the operating system base.

Installation and Support

This is a turnkey system; Avatar will coordinate all aspects of hardware acquisition and installation, as well as software installation and training. The environmental requirements of the system are modest, typical of the class of large minicomputers on which it runs. A raised floor is optional. Avatar recommends the use of surge suppressors on the power line.

System availability is quoted as 98%. Hardware maintenance is supplied by the selected hardware vendor(s); software support is provided by Avatar on a five day a week basis. Diagnostic software is available for all hardware components.

Source code for the applications, but not for the systems software, could be supplied to RLG. Avatar is willing to develop custom extensions to the system. Enhancements are made available at no cost to customers who maintain a software maintenance contract with Avatar.

System Operation

MIIS provides only some of the system control and utility functions commonly available with other operating systems. Avatar claims that no dedicated operations staff is required. We disagree: at least 1 FTE must be accountable for the well-being of the system, especially in regard to terminal network control and management.

System Administration

The system provides password-protected user identifiers. It provides a complete facility for assigning privilages to users, terminals, user groups, and terminal groups. It does not produce an "attempted breach" report to allow identification of repeated attempts to perform unauthorized functions.

Tailoring and profiling functions are supported to the required level.

Command and resource consumption logging can be enabled or disabled as required; some response time and resource consumption reporting facilities are provided. Correctable hardware errors are logged and can be reported as necessary; uncorrectable errors are not logged. Only disk utilization and file size reports are available to the system administrator; no other summary performance reports are produced.

Avatar estimates the required system administration staff as 2 FTE.

General User Interface

The system has as a design goal the ability to use any ASCII asynchronous terminal. It is possible to use a terminal that does not use cursor positioning commands, but the displays are not attractive. Avatar's terminal of choice is the Lear Siegler ADM-11. For applications requiring the ability to input and display the full MARC/ALA character set, they provide an HDS-108 with a custom character set PROM.

The command language is composed of English words, which can be abbreviated. Help screens are generally context-specific. For an inexperienced user, they can be invoked automatically. An installation can modify the help text.

Data Base Searching, Retrieval, and Maintenance

All published U.S. MARC formats are supported, except authorities.

An installation can specify the indexes to be maintained, and can specify the record fields and subfields to be included in each. Both word and phrase indexes are supported; a stop list is maintained for building word indexes, but stop words are not removed from search requests. Boolean searching is supported, as is implicit right truncation. Search results can be further restricted on the basis of non-indexed data elements. An index browsing function is provided.

All transactions to modify the data base are written to a queue, and then applied to the data base by a background program; index updates are done at the time the transaction is applied. There is no support for batch-mode data base maintenance. There is no recovery mechanism. If the system crashes while doing an index B-tree block split, for example, the only way to restore the data base to a consistent state is to request Meditech to patch blocks; alternatively, the data base could be reloaded from the last full dump, thus losing any transactions applied since that dump. Avatar is developing software to maintain a transaction audit trail by copying transactions from the queue to a disk- or tape-resident journal file at some specified frequency. Such a journal file would enable recovery via transaction rollforward from the most recent data base dump.

On-Line Catalog

Patron on-line catalog access is supported. In general, a user starts a search by selecting a search type (e.g., author, title, subject, call no., ISBN, corporate name). The system then displays the index so that it can be browsed for the desired term. Once the term has been selected, the system displays the associated title(s). Alternatively, the term may be entered directly. The on-line catalog shows the current status of all items. Searching may be limited to a specific branch or collection if this is desired.

Cataloging

Cataloging is partially supported, and the on-line entry of cataloging data is at a primitive level, suitable only for limited use.

An interface is provided to transfer OCLC records directly into the on-line catalog.

Authorities

An authority file is maintained as a component of the on-line catalog. The records are derived from authority-controlled tags in the bibliographic records; authority records may be edited, resulting in change to all linked bibliographic records, but they may not be entered directly. MARC authorities format is not used.

Acquisitions

An acquisitions function is not supported.

Serials Control

The serials control subsystem is still under development. It now supports check-in and bindery control to some extent; support for claiming is planned for sometime in 1984. Record access seems particularly flexible. Very little is provided in the way of reporting at the present time.

Circulation

The circulation subsystem appears to provide all necessary function.

Inter-Library Loan

There is no support for inter-library loan.

Vendor Qualifications

Avatar Systems, Inc. is a privately held corporation owned by its principal officers and a group of private investors in the DC area and Utah. It appears to have no more that 20 employees at present. The system is installed in several locations; present clients seem to be mostly technical and medical libraries. It is of interest that Carnegie-Mellon has picked Avatar as its system of choice.

Cost

No costed configurations were supplied.

3.2 Biblio-Techniques

Biblio-Techniques, Inc. (BT) is a small company based in Olympia, Washington. BT markets a vendor-supported "object code only" version of the software system owned and operated by the Washington Library Network (WLN). The system is called the Biblio-Techniques Library and Information System (BLIS).

Architecture

The architectural model to which BLIS conforms is the **single system with multiple applications**. BT may bow to requirements for a local area network interface, and would be interested in a joint development project with RLG in that arena.

Linkage Considerations

WLN is one of the participants in the CLR-funded Linked Systems Project, so the WLN system will have linkage capability in place by mid-1984. However the BLIS system is not identical in all respects with the WLN system, so additional work will be necessary to allow the WLN linkage software to be used.

BT has verbally committed to undertake the necessary development to link BLIS to the RLIN Data Resource Node. A memorandum of understanding to that effect is under negotiation at this writing. As a part of the effort BT would like to investigate the possibility of implementing all of the linkage support in the host, and/or using IBM-based interfaces instead of the DG mini front-end used by WLN and the Library of Congress.

Hardware Description

BLIS runs on IBM-compatible mainframes. The disk subsystem consists of CDC plug-compatible 3350's, although any PCM disk device that will connect to the mainframe may be used. The communications network is based on the NCR Comten front-end; the protocol is byte-synchronous over polled, multi-point lines. The primary terminal is the Telex 476, with the MARC/ALA character set. The Telex is compatible with the IBM 3270. ASCII terminals are also supported through the use of a protocol converter.

Underlying Software

BLIS uses the ADABAS data base management system, marketed in the United States by Software AG of North America, Inc. ADABAS is used at a relatively low level: as a record manager. Other services provided by ADABAS include BT/ET logic to guarantee the atomicity of multi-record transactions, and data base restart/recovery.

The major departure of BLIS from WLN is in its use of the Software AG teleprocessing monitor COM-PLETE, rather than IBM's CICS. COM-PLETE is said to be more robust than CICS and significantly easier to install and maintain. It makes use of the memory protection hardware of the 370, guaranteeing that no application can destroy another, or harm the operating system. Up to 15 threads of execution may be configured. Application code and all but 175KB of the COM-PLETE nucleus can be paged or not as desired; user memory is swapped. The system supports full conversational applications (i.e. is context-retaining); it imposes no limit on program size beyond that implied by the machine architecture. Monitor overhead is said to be equivalent to or lower than that of CICS. COM-PLETE includes a terminal mapping utility that makes it possible to design custom input screens easily. BT has put this to good use in the implementation of an on-line catalog interface. The switch to COM-PLETE was straightforward because the WLN software was designed with virtual monitor and data base system interfaces.

The operating system environment can be IBM's DOS/VSE, VS1, or MVS, depending upon the size and complexity of the system. Changes from one of these to another can be done without modification to the application. Furthermore, ADABAS can fully utilize multiprocessor IBM mainframes and supports a channel-to-channel connection between a CPU running application code and a back-end doing data base operations. These features permit upgrades to a very large maximum configuration.

The application implementation languages are PL/I (70%) and 370 Assembler Language (30%). Of note is a report writer called NATURAL, which is marketed by Software AG to be used in conjuction with ADABAS. This tool allows system implementers as well as users to easily define reports to suit ad hoc needs or for use on a continuing basis.

Installation and Support

BLIS can be categorized as a turnkey system -- one which can be installed quickly using off-the-shelf hardware without the need for user programming.

The environmental requirements include a raised floor with an under-floor plenum capable of dealing with 100,000 BTU/hr to maintain an ambient temperature of 70-76 degrees F. Ideally, the data center hardware should be attached to a 75 KVA power distribution center, which in turn requires connection to a 480-volt 3-phase power line. At least 1000 square feet of floor space should be provided for the configurations quoted in the RFI response.

Hardware maintenance is performed by the hardware vendor(s). Contracts are available on a 5-day prime-shift or on a 7-day 24-hour basis, depending on local requirements. Since several vendors are potentially involved, BT plays a coordination role in dealing with ambiguous problems.

All software is maintained remotely by BT for a monthly fee via a teleprocesing link. New software releases are provided and installed periodically at no charge beyond the monthly fee. Problem resolution is generally handled remotely over the link; priorities are assigned according to the level of severity.

On-site training is provided which covers the complete range of system function. A certain amount of self-training capability exists in the software as well. The user documentation appears complete, given the manuals now in existence and those currently in preparation.

It is too early to directly assess the availability and reliability characteristics of BLIS, since the first installation is just beginning production. Nonetheless, the record of WLN over the past several years together with substantial experience with the hardware technology leads us to belive that the availability in a well-run installation can be made to exceed 98%.

System Operation

The console support from the operating system conforms to the requirements in the RFI, as does support for terminal and data base control. A point of contention between BT and the site visit team centered around staffing levels for the operation of the machine. BT claimed that no dedicated staff is required. The site visit team feels that 2-3 FTE are necessary at a minimum if a dedicated system is used. Operator documentation and procedures are supplied with the system.

System Administration

The functions which deal with system security form a complete set, except for the lack of "attempted breach" reporting. System tailoring and profiling are supported at the required level, but BT staff must play an intermediary role in making the changes. Logging is performed on commands, with capture of both response time and resource consumption, but selective logging is not supported. Reporting is done on response only. Since the data is available, BT should implement the appropriate reporting to allow a fuller analysis of resource requirements, particularly in conjunction with use of the on-line catalog. Aggregate CPU and I/O activity should also be reported over intervals of a production day.

General User Interface

Both staff and patron terminals are full-face block-mode CRTs. The terminal, the printer, and the application software support superimposed diacritics. Input is performed by entering data in unprotected areas on the screen, with errors being flagged on a re-prompted screen.

Commands are common words, and can be abbreviated. The "help" functions are well-designed, being context-specific and oriented toward directing the user toward correction of his/her immediate problem. The help text can be easily modified by the library.

Data Base Searching, Retrieval, and Maintenance

All U.S. MARC formats are supported except for Maps, Archives and Manuscripts, and Authorities. The authorities subsystem now uses a non-MARC format, and is being upgraded to standard with a completion targeted for 1984. Support for Music has just been added.

Indexing may not be specified by the library. Besides the usual "number" indexes, subject and personal name are supported as phrase indexes; title and corporate/conference name are supported as word indexes. The system supports full Boolean capability. It does not support generalized browsing (authorities and call number browsing are supported), nor result limiting on date, country of publication, etc. The latter capability is planned for 1984.

When a bibliographic or authority record is added to the database, the old version is locked to prevent any further updates and the new version is placed in a queue. This queue is processed by a background task whose execution can be scheduled to suit installation needs. The new version of a record cannot be seen by other users until the update run completes; however a record can be seen using a special set of commands that access the queue directly.

Data bases may be loaded using the on-line background task described above (1000 records/hour), or by a more efficient offline process (up to 10,000 records/hour). During batch loading, editing is done at the tag level. Subfield-level editing is only performed on records entering the system on-line.

Online Catalog

The on-line catalog has a public access interface which is highly flexible and extremely well conceived. The WLN command mode (similar to RLIN) is supported, and is overlayed with a multi-level menu mode which is tailorable to some extent by the library staff. In addition to the provision of extensive "help" facilities, the menu mode gives the user a choice of several options in the construction of a search ranging from a step-by-step selection approach to direct command entry.

Searches are filtered through the authority file (in fact, headings reside only in that file, not in bibliographic records), and search re-direction on the basis of retrieved headings is accomplished without re-keying. An analogous -- and novel -- feature allows the user to re-direct a search using a tracing in a retrieved record as the next search argument. Heading indexes are constructed such that not only phrase and word searching are supported, but also subdivision and explicit non-subfield-a searches. Furthermore, sufficient data is carried in the indexes to provide substantially more assistance to patrons before incurring the overhead of needless large searches. These extensions could be added without extensive change to the underlying data base structure.

Cataloging and Authorities

The cataloging function conforms to the requirements in the RFI, with the notable exception of derivation of a new record from an existing one. WLN plans to develop this capability, but a date for completion has not been set.

Authority control is integral to the cataloging process, in that all incoming bibliographic headings are matched against the authority file, with unique headings being automatically added to that file and reported for review. While creating a bibliographic record, the operator can search the authority file and transfer a heading found there to the bibliographic record being constructed. Both bibliographic records and authority records are entered online in a full-screen mode with full on-line edits. Various levels of review are provided for different record categories, and a comparison mechanism highlights the differences between versions of a given record.

Acquisitions

The acquisitions system provides virtually complete procurement and fiscal capabilities, and is reasonably flexible. For example, different commands make possible either the combined processing of receiving and invoicing, or the separation of those functions. Screens are well designed though numerous; transition between them seemed more logical and less frequent than in the GEAC system.

Serials Control

No serials control capability presently exists in BLIS, and development will not begin until at least 1985 given the current development committments. BT expressed a willingness to implement a system if RLG central staff could provide the specifications, which could provide a serials capability sooner than would otherwise be the case.

Circulation

BT is under obligation to provide a circulation control subsystem by late 1984. The RFI response indicated that with a very few exceptions (most notably a booking function for special equipment and facilities) the new system meets the requirements validated in the member survey.

Vendor Qualifications

The organization numbers less than ten staff members, albeit highly qualified and motivated ones. The small size of the staff is partially mitigated by a shared development agreement with WLN. The BT income statement shows a small loss during the past twelve months, and the balance sheet could look healthier. BLIS is currently being installed at the University of British Columbia, and the University of California, San Diego. Columbia University recently issued a letter of intent to acquire the system.

Cost

Two configurations have been designed -- one for an average library and one for a large library, as defined in the RFI. Components and costs are given below for the two cases.

Average Library

Data Center Hardware	$448,000
Software Installation, Documentation, Training	17,500
Software License (Software AG components included)	225,000
TOTAL INITIAL COST	$690,500
Software Maintenance (Annual)	39,600
Hardware Maintenance (Annual)	40,740
TOTAL ANNUAL COST	$80,340

Large Library

Data Center Hardware	$523,000
Software Installation, Documentation, Training	17,500
Software License (Software AG components included)	225,000
TOTAL INITIAL COST	$765,500
Software Maintenance (Annual)	39,600
Hardware Maintenance (Annual)	47,540
TOTAL ANNUAL COST	$80,340

3.3 Carlyle

Carlyle systems is a small company in Berkeley, CA. It was formed in June, 1981 to design, develop, and market turnkey library systems. Its first products are now under development, and the first installations are targeted for early 1984.

Architecture

In terms of the architectural models described in the RFI, the Carlyle system is a single system with multiple applications, even though in general it is built from multiple processors.

Linkage Considerations

The Carlyle systems, when complete, will comply with ISO Standard 7498: Open Systems Interconnection - Basic Reference Model and related protocol standards through level 4 (the transport level). Carlyle has stated that they are prepared to cooperate with efforts to implement levels 5 through 7.

Hardware Description

The system is based on use of Intel iDMS data base processors for data storage, with additional processors constructed from off-the-shelf Intel Multibus components. The iDMS itself is a Multibus system containing an 8086-based processor board, an 8089-based SMDI disk controller to which up to four disk drives can be attached, and an Intel Ethernet interface. The overall design involves a number of communicating processes, as shown below:

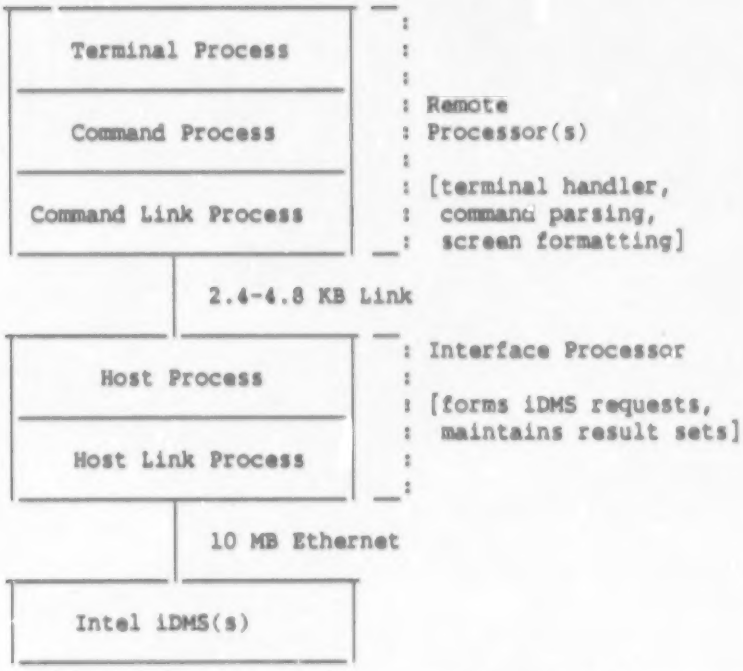

The rationale for separating the interface and remote processors is that in a large library or library system a number of remote processors in different locations could serve as terminal concentrators, communicating with a single interface processor and data base in a central location. In a small configuration the functions of the interface and remote processors could be combined in a single user processor. In very large configurations, multiple interface processors could be used. The largest remote processor is said to support 22 terminals; the interface processor is said to handle 50 terminals connected to multiple remote processors.

Carlyle plans to use Televideo 914 terminals, Century 300 MB disk drives, and Siemens ink jet printers. The designers of the system felt that a tape drive is unnecessary; we disagree strongly.

Underlying Software

The application software will be written in a combination of assembly language, PL/M (an Intel language), and Pascal. They intend to use the Intel RMX-86 real-time operating system. RMX-86 has been totally rewritten; its predecessors have a reputation for ponderousness. Release 5 of RMX-86 is a multiuser system and is presumably what they will use.

The iDMS provides for the definition of relational data bases with limited numbers of fixed length, singly occurring fields making up keyed files. The device performs relational algebra operations, provides BT/ET logic for transaction atomicity, allows the invocation of predefined named "macro procedures", and provides access to sequential files of unstructured data by relative byte address. The limitations of this structure will be overcome by storing bibliographic and other records in a MARC-like format in the byte-sequential files, and referencing them from the keyed files, which will serve as indexes. Thus, the iDMS will serve as an intelligent record manager. When a Boolean search involves only a single index, the entire request will be passed to the iDMS for processing. The interface processor will handle Boolean operations that involve multiple indexes, because any two indexes might not reside on the same iDMS.

Installation and Support

Since the development is not yet complete, reliability statistics are not available. The design goal is to provide 99% uptime. Hardware maintenance will be handled by the individual hardware vendors, or by a third-party maintenance firm. Diagnosis will be performed by Carlyle through a telecommunications link. The link will also be used to check the system on a daily basis.

One week of free training will be provided on-site. The adequacy of the user and operations documentation is unknown, since it was incomplete at the time of the site visit.

System Operation

The RFI response indicated that operations support would exist, but it seems incomplete without capabilities to initiate batch work or utility functions. Data base control functions must exist, but their speed of operation was not given.

System Administration

Privileges in this system will consist of predefined sets: one for staff and one for patrons. These sets will be modifiable by the library. Tailoring capability will allow for defining and altering existing display screens, and for the alteration of help text and printed products. Support will exist for the addition of new terminals, data base storage devices, and new elements to a file of the data base. Profiles will be definable to include default display format, data bases and files to be searched, etc.

Command logging and reporting will be supported, but not with respect to resource consumption. The reporting of CPU and I/O path use over time will likewise not be supported.

General User Interface

Staff terminals will be IBM PC's, configured with a graphics display and a hard disk. The staff user interface will be form fill-in with protected and unprotected areas on the screen. Wand and laser readers will be supported for circulation use also. The patron terminal will be a CRT operating in character mode, with optional attached printer.

Help information will always be available, and will be automatically displayed in some circumstances. The help will be context-specific. Local staff may modify the help text through the use of a utility program.

Data Base Searching, Retrieval, and Maintenance

The library will be able to specify the indexes to be created, as well as the elements and subfields contributing to each. Word indexes will be supported, but not phrase indexes. This is a drawback in the instance where, say, a title is known precisely and could be specified to allow for a more economical search. Boolean search expressions will be supported. Results will be reduceable by specifying non-indexed elements, but inequality operators will not be supported.

Indexes will be updated in real time for circulation data; the modification of bibliographic data will be reflected through overnight processing.

On-Line Catalog

In looking at a demonstration system, we find that the catalog will have most of the requesite functions, and the tutorial capability will be integrated into the overall system. Valid headings and cross-referencese will display with bibliographic data, and automatic switching will be performed if a SEE reference is retrieved.

Cataloging and Authorities

Cataloging functions will be restricted to the maintenance of records already in the data base. The RFI response indicated that a limited number of machine edits will be performed on modified records. A batch listing capability for bibliographic records will be supported.

The authority capability likewise will not allow the addition of new authority records to the file directly; records will be added as a consequence of new headings entering the system through bibliographic records. Machine edits and related record checks will be limited in scope.

Acquisitions and Serials Control

Acquisitions and Serials Control will not be supported by Carlyle. Discussions have been held between Carlyle and Innovative Interfaces pursuant to a linkage between those two systems.

Circulation

Circulation, when supported, will provide most capabilities cited in the RFI, except for Special Materials Booking.

Interlibrary Loan

Carlyle will not offer an interlibrary loan function.

Vendor Qualifications

The proposed product is still under development, and hence there are no current installations.

Carlyle Systems is a California limited partnership, 5.75% owned by individual limited partners, 2% owned by a venture capital firm, and the remainder by Carlyle Systems, Inc. The latter is wholly owned by Stephen R. Salmon.

The staff consists of the President, and senior officers for product development, marketing, and finance. Technical staff work for the company as independent contractors.

3.4 CLSI

CL Systems, Inc. is based in Newtonville, Massachusetts, and markets turnkey library systems. It has the largest installed base of any vendor of library systems.

Architecture

The current CLSI system is best described as a single system with multiple applications. Larger configurations are actually implemented with multiple processors, interconnected via high-speed links. Development plans call for replacing the current multiprocessor linkage hardware with LAN technology; when that is done, the system could well evolve in the direction of an integrated collection of application servers on a local area network.

Linkage Considerations

CLSI has agreed to work with RLG in the development of the facilities necessary to link a CLSI system at an RLG member site with the RLIN data resource node.

Hardware Description

The CLSI system runs on DEC PDP-11 hardware; currently, PDP-11/24 and PDP-11/44 processors are used. Multiple processor configurations are interconnected via DA-11 Unibus links.

The LIBS 100 multiprocessor system may be configured with up to 5 processors, one of which is the "data base processor," the rest of which serve as terminal processors. One of the terminal processors may serve as a hot spare for the data base processor if it is configured with the necessary links to the other terminal processors and to the disk drives on which the data base is resident. All terminals are connected to the terminal processors via an intelligent port selector so that they are easily reassigned if one of the terminal processors is pressed into service as a data base processor.

Underlying Software

The operating system employed by CLSI is of their own design and implementation. Presumably, it is streamlined and optimized to support their applications efficiently. Applications are written in a dialect of PDP-11 assembly language (dubbed "FLIRT" by CLSI).

Installation and Support

Systems are completely configured and tested by CLSI prior to shipment to a customer. Thus, customer-site installation is relatively simple and the system may be characterized as a turnkey system. Environmental requirements are modest. A raised floor is not required, but is considered desirable by CLSI, especially for larger configurations. Use of a line conditioning transformer is recommended for environments that experience problems with electrical noise on supply lines.

Average uptime for installed systems was quoted at 98.4% with a very small standard deviation. For a multiprocessor configuration, it was quoted at 99%. Diagnostic software is provided for all supported hardware. Both hardware and software maintenance are performed by CLSI personnel. All systems are supplied with an auto-answer modem or are being retrofitted with one to facilitate remote problem identification and maintenance. This support is offered on a seven day a week basis.

Source code is not supplied to customers, but could be made available to RLG under special arrangement. CLSI makes a new general release of its software to all customers once a year, with interim releases to selected customers as needed to meet contractual obligations or to repair especially onerous problems. They will also develop and maintain special software for a specific customer.

Documentation, training, and conversion support appear to be adequate.

System Operation

All common system control functions are available to the operator of a CLSI system. Isolation and patching of terminal lines are supported. The daily data base backup procedure seems complex, however.

CLSI estimates the required operations staff as 2 to 3 FTE.

System Administration

The system provides a complete facility for assigning privilages to users, terminals, user groups, terminal groups. It does not produce an "attempted breach" report to allow identification of repeated attempts to perform unauthorized functions.

Tailoring and profiling functions are supported to the required level, except that profiles can only be assigned to individual users and terminals, and not to user or terminal groups. Furthermore, new display and input screens cannot be defined by the installation (at least without the participation of CLSI), although existing ones can be modified.

Command and resource consumption logging can be enabled or disabled as required; extensive response time and resource consumption reporting facilities are provided. Both correctable and uncorrectable hardware errors are logged and can be reported as necessary. All required summary performance reports can be generated by the system based on its log data.

CLSI estimates the required system administration staff as .5 to 1 FTE.

General User Interface

Both staff and patron terminals are character-mode CRTs. The applications interface is screen-oriented; option choice may be made either with function keys, with a light pen, or via a touch-sensitive screen. The CLSI terminals do not support the full MARC/ALA character set, with the exception of an Apple-based MARC editing terminal. Data input is performed by filling in unprotected fields of a preformatted input screen.

The command language is composed of English words, which can be abbreviated; it seems quite complex. Help screens can be invoked whenever required and are generally context-specific. An installation can modify the help text.

Data Base Searching, Retrieval, and Maintenance

Currently, the system supports records that are considerably simpler than those specified by the U.S. MARC standard. All material types are supported in the simplified form, however. The mid-1984 software release will provide for the storage, retrieval, and maintenance of full MARC records.

An installation can specify the indexes to be maintained, and can specify the record fields and subfields to be included in each. Only phrase indexes are supported; the mid-1985 release will support word indexing. Boolean searching is supported, as is both explicit and implicit right truncation. Proximity searching is planned for the mid-1985 release. Search results can be further restricted on the basis of non-indexed data elements. An index browsing function is provided.

According to the RFI response, some indexes are updated on-line, the rest off-line. Record additions, replacements, and deletions may be done in batch mode. The recovery mechanism presently employed is daily backup of the data base to removable disk packs; support for streaming tape drives is planned. Support for before-and after-image logging of disk updates will be added in the mid-1984 release.

On-Line Catalog

An on-line catalog is available now, but it is undergoing major change and enhancement; the new version is promised for mid-1984.

Cataloging

CLSI does not provide a real cataloging system at the present time. The principal means of entering records into the online catalog is batch loading of records cataloged on a bibliographic utility, usually OCLC. Support for the on-line creation and modification of MARC records is under development; it will make use of the Apple-based MARC editing terminal.

Authorities

The present system does not include authority control. Implementation of an authority control system is planned for mid-1984.

Acquisitions

Acquisitions system functionality is relatively complete, although claiming and a management reporting facility will not be supported until the mid-1984 release.

In general, a user interacts with the acquisitions system through a long series of functional steps, whose sequence is rigidly defined. It is not possible to make detours, e.g. to look at vendor or fund information, without backtracking step by step through the levels of the process, or popping all the way to the topmost level. Most of the other acquisitions systems reviewed allow such "lateral" movement.

The only access to a record in the receiving or invoicing process is by vendor.

Serials Control

Serials control is not now supported, but is under development, with mid-1985 availability. As an interim measure, CLSI is investigating serials control systems they could recommend from other vendors.

Circulation

Circulation is the original CLSI application; the vast majority of installed systems are circulation systems. Functionality is complete, for the most part. It is currently not possible to alter the order of the hold queue, but this facility will be available in mid-1984. Patron record update requires as many individually prompted transactions as there are fields to be processed.

Inter-Library Loan

Inter-library loan is supported by the printing and mailing of printed forms. The user interface to this function is not screen-oriented; instead, data elements are entered singly in response to prompts. A more extensive system based on electronic transmission of requests is planned for release in mid-1984.

Vendor Qualifications

CL Systems Inc., a privately held corporation, is said to have completed the fiscal year ending June 30, 1983 with before tax profits of $1.4 million on gross revenues of $18 million. This represents substantial improvement for a company that two years ago was in serious difficulty and underwent a major reorganization. There are more than 100 employees. At the time of our site visit, there were 660 libraries using CLSI systems (over 200 installed systems, some of them shared).

Cost

No costed configurations were supplied with the CLSI response.

3.5 GEAC

GEAC, Ltd. is based in Toronto, and markets turnkey systems for libraries and financial institutions. It has a very sizable number of library systems installed in the United States.

Architecture

GEAC does not conform strictly to any of the models hypothesized in the RFI. Currently, the GEAC configuration can be described as multiple applications running on one or more homogeneous discrete systems. Up to two systems can be linked, either using shared memory or via an RS232 interface and a virtual terminal protocol. In the future, more than two machines will be able to link through attachment to a common bus. At the present time, relationships among files on different machines are maintained using batch data transfer. Some consideration is being given to the use of a local area network for subsystem interconnection in the future.

Linkage Considerations

GEAC has agreed to work with RLG in development of an intersystem linkage between an RLG member site and the RLIN data resource node. A memorandum of understanding to that effect is under negotiation at this writing. While GEAC has not participated in the CLR-funded system interconnection project in the U.S., they are working on its Canadian counterpart, and are well aware of what they will be committing to. They are strongly committed to the ISO Open System Inteconnection Reference Model, and feel positively toward standards activities now under way.

Hardware Description

GEAC currently manufactures their own computers. While off-the-shelf components are used, the architecture can nonetheless be considered unique. The technology employed is MSI TTL, including the 2901 bit-slice ALU chip.

The GEAC hardware product line ranges from a small desk-top unit (the 2000) to a multiple-processor unit (the 8000). It is this latter unit that is of primary interest; therefore the rest of this section will focus on the 8000.

The 8000 contains four processors, each of which is dedicated to a special function: disk I/O, tape I/O, communications I/O, and application processing. Each processor has its own local memory, plus the ability to access a global memory. The system can support in a hardware sense two megabytes of global memory, and two megabytes of local memory for each processor. In practice, however, the communications and disk processors are seldom configured with

more than 128 kilobytes each, whereas the shared memory and application processor frequently have 500 kilobytes each. It is possible to configure a dual 8000 system with a common shared memory if more processing capacity is desired.

The communications processor manages all attached terminal devices, as well as high-speed printers. It can support up to 64 ports and has sufficient bandwidth to handle about 50 9600-bit/sec lines. The protocol used with the current library terminals is X3.28 (polled, block mode). X.25 Datapac support is also offered.

The disk processor manages all attached disk storage and provides a cache. Storage Module Device Interface (SMDI) devices are supported; the present product line includes 160, 300, and 675 megabyte devices. The disks are OEM devices not manufactured by GEAC. There is a single data path to the disk string at the present time.

An application program is divided into code and data segments, each up to 64 KW. Typically, the code segment is loaded into local memory and the data segment is loaded into shared memory. The local memory is virtual, but production systems are executed virtual=real to eliminate the overhead inherent in paging.

The GEAC 9000, currently in development, is a generalization and extension of the firm's current multiprocessor architecture. A single 9000 system may have up to 8 processors, connected to up to 64 megabytes of memory via a crossbar switch. The bit-slice processors are 16 bits wide internally, but can address all 64 MB of memory without context switching. The individual processors are expected to have 60% greater throughput than the aggregate throughput of an 8000 system. The architecture includes multiple register sets for fast context switching. Devices of a particular type can be attached to any processor, unlike the present system in which each processor is functionally specialized. Furthermore, multiple 9000 systems may be interconnected via a high-speed bus. The instruction set is upward-compatible with the current product line, so that existing applications can be transported easily.

Underlying Software

The operating system employed by GEAC is one of their own design and implementation. Its design is optimized toward transaction processing with low overhead. The implementation language for both system and applications is ZOPL (also unique to GEAC). A record manager provides access to records given their keys. File and data base functions are implemented in the application. Some consideration is being given toward moving to a standard implementation language such as "C", but the cost of conversion will probably limit the use of that language to new projects, if it is chosen.

Installation and Support

GEAC can be categorized as a turnkey system -- one which can be installed quickly using off-the-shelf hardware and which requires no user programming prior to installation.

The environmental requirements of the system are modest, with the "large library" configuration requiring air conditioning to remove 40,000 BTU/hour. A raised floor (100 - 150 sq.ft.) is recommended. The use of a power protection unit is optional. Power requirements are 60 A, 208/220 VAC, single phase.

MTBF for the system as a whole is quoted at 4000 hours under normal conditions, but less (unspecified) during the three months following a new software release. Sizing formulae are available for circulation control, but not for the on-line catalog. Hardware maintenance is performed by GEAC. Diagnostic software is provided for all supported hardware. Changes in the MARC formats are handled by modifying data tables that are externally accessible. Internal documentation and source code is available to RLG or to any using member institutions. Bug fixes and functional enhancements are provided in new releases of the software without charge. Documentation and training support appear to be adequate.

System Operation

All system control functions generally available to an operator are present in the GEAC system. Isolation and patching of terminal lines are supported, but network diagnostic capabilities await the existence of a GEAC local area network. Data base recovery functions are adequate. The requisite console utility functions are all present except for the capability to produce a formatted memory dump. Operational staffing is on the order of 2-3 FTE. Operator documentation and procedures for often-performed tasks are available.

System Administration

The system has a relatively complete set of functions for system security and privileges. The system does not produce an "attempted breach" report to allow identification of repeated attempts to perform unauthorized functions. Tailoring and profiling functions are supported to the required level, except that new display and input screens cannot be defined by the library. The library can, however, modify preexisting ones. Command logging is always in effect, but no resource statistics are kept. This tool is used exclusively to monitor response time. Uncorrectable hardware errors are not logged for later reporting, but correctable ones are. Summary reports on system CPU utilization and disk space are available on demand.

General User Interface

The primary support for terminals, both staff and patron, is the full-face block-mode CRT. The full MARC/ALA character set is supported and diacritics are superimposed on the characters they modify. Input is performed by filling in unprotected areas of a preformatted input screen. Errors are flagged on a re-prompted screen.

The command language is composed of English words, with the capability to abbreviate. Help screens are available whenever desired. They do not in general indicate how the user can correct his/her immediate problem, however. A using library has the capability of modifying the help text using one of several methods.

Data Base Searching, Retrieval, and Maintenance

All U.S. MARC formats are supported. A using library may designate the indexes to be maintained by the system, and can specify the fields and subfields to be included in each. In general, indexes are phrase indexes, except for the "key" index in the online catalog composed of words from author, title, and subject fields. Boolean search expressions are currently unsupported, but this is scheduled for implementation in early 1984. Results can be limited using publication date, country of publication, etc., together with relational operators. Also to be added in 1984 will be adjacency, improved search status display, better user guidance for searching, and support for search result identifiers.

Index updating is performed partially in real time, and partially in an overnight process. Batch maintenance for adds, replacements, and deletions is completely supported.

On-Line Catalog

The on-line catalog is currently running in pilot mode at New York University. Access is via author phrase, title phrase, subject phrase, author/title phrase (requiring two screens for input), numeric elements, or "key" -- an author, title, or subject word. As indicated above, Boolean expressions are not yet available, but will be in 1984. Integration of the on-line catalog with data created in other subsystems appears quite complete.

The tutorial mode explains the different access points, though it takes several screens to do so fully. The adventurous user can graduate to command mode in a fairly short time, and help text is available at any point in a session.

An author or title search which does not match any index entry presents the user with a window on the appropriate index; after a single-record result the user can specifically invoke such a window relative to the current search argument. Levels of record display range from multi-record 'brief' to single record partial to a fully tagged record. Although the user can backtrack through a sequence of earlier screens, the context in which a command operates never spans screens, obviating the possibility of refering to a record on a previous page.

Cataloging

The closest thing to a cataloging function is provided by the MARC Record Management system (MRMS), a very basic capability to edit batch-loaded MARC records and, for experienced operators, to create new records at the terminal, one field per screen. GEAC has not heretofore seen a market for a cataloging subsystem owing to the services provided by the utilities. They have, however, concluded an agreement with the Smithsonian Institution to provide a full-function cataloging system by mid-1985. Presumably, this would be implemented in a sufficiently general way as to be useable by other GEAC users who desire this support.

Authorities

An authorities subsystem is now being designed for NYU with an expected completion in mid-1984. From the RFI response, it is seen that most of the functional requisites will be met by this with the exception of related record edits on new authority records (e.g. non-duplication of the same 1xx field in another record). Our impression was that the requirements have not been completely analyzed.

Acquisitions

According to the RFI response, the support for Acquisitions is functionally complete and in the first phase of installation and operations testing. In demonstration, however, it appeared somewhat cumbersome to use. Straightforward transactions such as generation of basic bibliographic and order data to print a purchase order required a large number of screen transitions, each screen being used for a discrete sub-unit of the process such as vendor selection, fund allocation, etc.

Serials Control

It appeared from the RFI response that a serials check-in system is now supported, but it does not support control aspects such as prediction, handling, volume closeout, etc. However, it will be upgraded to provide full function by January, 1984.

Circulation

Currently, GEAC is best known for its circulation support. Aside from the inability to alter the order of a hold queue and to have the reason for blockage displayed if the system doesn't accept a chargeout, the circulation functionality appears to be complete given the RFI response.

Inter-Library Loan

GEAC supports an inter-library loan function, but messages can only be sent and received from other GEAC systems at this time.

Vendor Qualifications

GEAC, Ltd., publicly owned and listed on the Toronto Stock Exchange, had revenues of $48 million in fiscal 1983, with a net income of $3.7 million. Its cash flow position is conservative, and its long-term debt is nil. The Canadian government has in the past issued grants to GEAC for product development. The company has about 600 employees, with about one-third devoted to the development and support of library systems.

Cost

Three configurations have been designed -- one for each of the cases defined in the RFI. Components and costs are given below for the three cases:

Small Library

(1) CPU, GEAC 6000, 512 KB, 16 port MUX package	$90,850
(3) 128KB memory board	11,460
(2) disk drives, 332 MB	40,000
(1) tape drive, 100 IPS	6,800
(1) printer, 300 LPM	6,400
(1) dial-up modem	1,230
GEAC Library Information System (software for applications described above)	100,000
Installation and Conversion	9,000
Software Maintenance	10,800
Hardware Maintenance	16,080

Average Library

(1) CPU, 8000, 640 KB, 16 port mux package	$116,695
(3) 128 KB memory board	11,460
(3) disk drive, 332 MB	60,000
(1) tape drive, 125 IPS	13,200
(1) printer, 300 LPM	6,400
(1) dial-up modem	1,230
GEAC Library Information System	100,000
Installation and Conversion	9,000
Software Maintenance	10,800
Hardware Maintenance	21,324

Large Library

(1) CPU, 8000, 640 KB, 16 port mux package	$116,695
(7) 128 KB memory board	26,740
(7) disk drive, 675 MB	255,500
(1) tape drive, 125 IPS	13,200
(1) printer, 300 LPM	6,400
(1) dial-up modem	1,230
Boolean Processor	60,000
GEAC Library Information System	100,000
Installation and Conversion	9,000
Software Maintenance	10,800
Hardware Maintenance	63,564

3.6 Innovative Interfaces

Innovative Interfaces, Inc. (III) is a small company located in Berkeley, CA. III designs, develops, and markets library systems and inter-system communication devices.

Architecture

The III system can best be described as a single system with multiple applications. It could form a component of a heterogenous set of discrete systems.

Linkage Considerations

III has stated its willingness to cooperate with RLG in the development of an inter-system linkage. While there seems to be no intrinsic reason why the III system could not communicate with the RLIN Data Resource Node, the site visit team had the impression that III had not studied the issues to the point of understanding the nature and extent of the work required.

Hardware Description

III runs on an S-100 multiprocessor system. The processor chips are Zilog Z80s. A single "master" processor resides on the bus with several "slave" processors and a disk controller. Each processor has 128K of bank-switched RAM. There is no shared memory.

The primary functions of the master processor are to control bus access, to load the slaves, and to control concurrent access to the data base. When a slave requires service, it interrupts the master and passes its request over the bus. With one exception, the slave processors run the application code, each supporting two terminals. The application processors perform data base reads directly, but rely on a single dedicated slave processor to preform all data base updates. Updates are queued for the data base update process in a transaction log file to which any application processor may append records.

Priam 120 MB winchester disks are used; up to four may be attached to a disk controller. Video tape is used at present as the backup medium; industry standard magnetic tape is being investigated as an alternative.

III is presently investigating alternative hardware bases; candidates include Intel's recently announced 186 processor and Convergent Technology's 8086/68000 system.

Underlying Software

The software base for the system is the UCSD p-System, with modifications for multiprocessor support. The applications development language is Pascal. III has developed a record manager to support multi-indexed files.

Installation and Support

III uses a communication link to remotely diagnose problems. The purchaser is expected to be able to carry out board swaps and related procedures, in consultation with III in Berkeley.

On-Line Catalog

III expects to bring an on-line catalog to market in early 1984. It will be targeted at small libraries, and will be a simple no-frills capability.

Acquisitions

The primary product at this time is a stand-alone full-function acquisitions system, which encompasses both procurement and fiscal processing. The system is easy to use, and seems to fit well into the usual workflow.

Serials Control

This product was still under development at the time of the site visit, but was demonstrable. It appears have good functionality. The screens are designed to emulate the appearence of a cardex file; some find this idea attractive, and some find it "quirky."

Costs

Although III did not include the requested cost information in the RFI response, we were given some prices during the site visit. Current list for a 3-processor, 2-terminal configuration is $55,000. An 8-processor 10-terminal system is $120,000. The average currently installed system cost about $70,000. Maintenance costs are 1% of the purchase price per month.

3.7 NOTIS

The NOTIS system has been designed and implemented in the library at Northwestern University (NU) over the past decade. Northwestern now markets the system and assists in installation and training. NOTIS is installed in several libaries world-wide, and is in the process of being installed in several research libraries in the United States.

Architecture

The architecture of the NOTIS system is that of a single system with multiple applications. Because NU is moving in the direction of a digital PBX, there are no plans to implement a LAN interface.

Linkage Considerations

NU has made a verbal committment in principle to implementing intersystem linkage as defined in the RFI. They have initiated discussions with potential granting agencies with the purpose of securing outside funding for that project.

Furthermore, the NOTIS Development Manager at NU performed the design on the LSI Application Level Protocol under contract to RLG; therefore NU staff have a detailed understanding of what is involved in such an effort. The SNI implementation at the Library of Congress will have considerable value to the project, since the hardware and software at LC are similar to that at NU.

Hardware Description

NOTIS runs on IBM-compatible mainframes. The possible disk configurations are those generally available to run in conjuction with those machines. Communication support exists for the IBM 3270 CRT terminal and for its look-alikes. In addition, ASCII devices are connected to the system using a protocol converter.

Underlying Software

NOTIS uses IBM's CICS as a transaction monitor. There are versions of NOTIS that will run either under IBM's DOS/VSE operating system, or under MVS. The VSAM access method is used as a record manager. The BTAM, VTAM, or TCAM access methods support terminal communications. The online protions of NOTIS are implemented in IBM 370 Assembler Language, and PL/1 is used to implement some batch programs.

The functions normally associated with a data base management system are spread throughout the application. Currently, no capability exists to guarantee transaction atomicity with respect to the data base.

The cataloging component of the system is implemented as a conversational CICS application: context is maintained between transmissions to the terminal. The online patron access task is pseudo-conversational: the task is terminated on transmission to the terminal, but some context is maintained in temporary storage by the application itself to simulate a conversational application.

Installation and Support

NOTIS is not a turnkey system. While the hardware configurations consist entirely of off-the-shelf components, the software is likely to require a certain amount of tailoring.

The environmental requirements include a non-carpeted floor and an air conditioning system capable of dealing with between 20,000 and 40,000 BTU/hour, depending on whether the configuration is for the "small", "medium" or "large" case. Power requirements vary between 10 and 20 KVA, 110/208 VAC, 1 and 3 phase. Machine room space requirements are approximately 300-400 sq. ft.

Hardware maintenance is performed by the hardware vendor(s). It is up to the library to determine problem response requirements and negotiate a contract accordingly. Contracts can be made for 5-day prime shift, and 7-day 24-hour responsiveness.

NOTIS software maintenance includes enhancements to the system and to the documentatio. Telephone assistance is provided in support of problem isolation and correction, but it is assumed that a technically competent individual will be present in the user's site. Each licensee is entitled to 15 FTE-days of installation support and training. Additional FTE-days will be billed to the customer.

Although statistics are not available for MTBF and MTTR, the measured ratio of uptime/scheduled time is 99.5% at NU, which is reliable by anyone's measure.

System Operation

System operation is simple and straightforward, and the required functions exist. No provision is made for network control: either diagnosis or hardware bypass. Data base control facilities are simple, and consist of full and differential dump/restore.

System Administration

The security functions are complete except for the provision of an "attempted-breach" report. A profiling capability does not exist. Tailoring operations may be performed, such as modifying an existing display screen or changing "help" text, but they must be performed by a programmer and require re-assembly of parts of the system.

Logging is performed for all patron commands issued. Response is recorded, but not resource consumption. Reporting consists of command frequency reports and session chronologies produced on demand. Other kinds of hardware resource monitoring (e.g. CPU utilization over time, I/O path utilization over time) is available as a priced software option from IBM.

An optional, general purpose report writer can be used to produce a wide variety of management reports without programmer involvement.

General User Interface

The primary terminals employed in NOTIS are block-mode CRT terminals. For staff performing cataloging functions, data entry is free-form: <tag><data><field separator>. This format accomodates variable-length data without recourse to programmable terminals. For other data, full screen forms with unprotected data areas are used. Patron access terminals may also be ASCII terminals, connected to the system using a protocol converter.

Familiar natural language words are used to form commands. "Help" is always an available patron option, but is not performed automatically by the system. Help screens are context-specific, and are designed to assist the user in solving the immediate problem. Library staff can modify the help text through the mediation of a programmer as described above.

NOTIS is designed to complement or enhance efficient work flow: input/update screens are cleanly presented, multi-screen transitions are minimized for any given function, and prompts and defaults reflect a real understanding of library practices.

Data Base Searching, Retrieval, and Maintenance

All published MARC formats are supported. The library can select which of the defined indexes it wishes to use. Phrase indexing with implicit right truncation is employed throughout; there are no current plans to support word indexing or Boolean searching. Records are limited in size to 7200 bytes owing to the requirement that a record must fit within a VSAM control interval. Changing the interval would require some programming changes to the batch system, and would have a negative effect on performence. It is estimated that 2-2.5 million bibliographic records per institution can be accomodated without system modification.

Of special interest is the MARC batch searching component. Technical processing staff can issue a search request on-line which will be stored and later passed against a sequential MARC file on tape. Hits are stored for use as copy. The searching operation can be performed as often as the library deems appropriate. Standing searches are also supported.

Online Catalog

The patron access component is well-constructed, with straightforward tutorial facilities and very simple syntax. Displays are generally well designed, and make use of half- and full-intensity to highlight significant data. The distinction between monographs and serials is not particularly intuitive, but to the regular user that would not be a problem. The principal shortcoming of the system has been the inability to move directly from bibliographic data to circulation status information; the release of a new, integrated circulation system will remove that infelicity in mid-1984.

The NOTIS on-line catalog does not provide tutorial capability, but the need for this is doubtful given its simplicity, its use of prompted menus for commands, and the availability of context-specific help screens.

Cataloging and Authorities

Mention has already been made of the use of free-form screens to enter and modify variable fields in a bibliographic record, and of the off-line MARC searching capability. With the exception of on-line validation of headings (see below) all edits are supported. A batch program gives the ability to make global changes.

While a separate authority file exists, the logical authority system includes both bibliographic and authority data, in that separate authority records are not usually created for headings that have no references and do not conflict with other headings. Records in the authority file are not yet linked to headings in bibliographic records, but a set of integrated Bibliographic/Authority indexes is under development for completion in mid-1984. Cross-references will be included. At that time, headings will be validated through the indexes. Conflicting headings will be suppressed from public view and reported to Cataloging.

Acquisitions

The acquisitions subsystem provides for order and receipt processing, and to a lesser extent the maintenance of fiscal data. Positive comment has already been made regarding the ease with which these functions can be performed. A complete fiscal package is planned for implementation by the first quarter of 1985.

Serials Control

A set of check-in functions now exist to record receipt at the item level, and to exercise some bindery control. Piece prediction, volume close-out, and automatic summary holdings update are not yet supported. Claim and renewal alert lists are generated, but not binding alert lists, supplier performence logs, management statistics, bindery slips, or material routing slips.

Circulation

A circulation component will be released in the first quarter of 1984. Support for reserve processing will be added in early 1985. The RFI response showed planned support for very few batch products beyond the usual patron notices; however the NOTIS staff intends to make use of a generalized report writer to produce the reports necessary to manage the circulation function.

Inter-Library Loan

Thought is being given to supporting an ILL function, but no date has been assigned to this project.

Vendor Qualifications

NOTIS has been installed in more than ten major libraries in the United States and abroad. NU currently employs 6 FTE in software development, and 2 FTE in marketing support. An additional number (2-5 FTE) will be added over the next twelve months. The libraries at NU, with the support of the administration, have made a strong committment to the continued enhancement and support of the NOTIS system.

Cost

Three configurations have been designed -- one for each of the cases defined in the RFI. Components and costs are given below for the three cases:

Small Library

(1) IBM 4321 CPU, 1 MB	$60,342
(1) System Console	2,988
(1) Maintenance Terminal	1,462
(2) disk drives, 571 MB	52,768
(2) tape drives	18,279
(1) line printer, 650 LPM	28,794
IBM operating system license	17,000
NOTIS software license	58,000 [1]
Report writer software license	8,500
IBM software rental	3,540
NOTIS maintenance	7,000
Report writer maintenance	595
Hardware maintenance	17,886

---

1. First year maintenance must be purchased simultaneously.

Average Library

(1) IBM 4331 CPU, 2 MB	$91,090
(1) System Console	2,988
(1) Maintenance Terminal	1,462
(3) disk drives, 571 MB	75,378
(2) tape drives	18,279
(1) line printer, 650 LPM	28,794
IBM operating system license	17,000
NOTIS software license	58,000 [2]
Report writer software license	8,500
IBM software rental	3,540
NOTIS maintenance	7,000
Report writer maintenance	595
Hardware maintenance	20,466

---

2. First year's maintenance must be purchased simultaneously.

Large Library

(1) IBM 4361 M5 CPU, 8 MB	$219,360
(1) System Console	2,988
(1) Maintenance terminal	1,462
(4) disk drives, 730 MB	105,536
(2) tape drives	42,755
(1) line printer, 650 LPM	28,794
(1) communications controller	42,449
IBM operating system license	17,000
NOTIS software license	58,000 [3]
Report writer software license	8,500
IBM software rental	3,540
NOTIS maintenance	7,000
Report writer maintenance	595
Hardware maintenance	29,748

---

3. First year's maintenance must be purchased simultaneously.

3.8 RLIN

RLG member institutions currently make use of the RLIN system, running at Stanford CA, for shared cataloging, acquisitions procurement, reference searching, interlibrary loan, and accessing subject-specific data bases.

Architecture

RLIN is a single system with multiple applications.

Linkage Considerations

RLG is one of the participants (along with WLN and the Library of Congress) in the Linked Systems Project. It is therefore the case that RLIN will be capable of linking with these systems in mid-1984. It will also be able to link with other systems that obey the protocol conventions, including another copy of itself. Incremental development to support linkage is therefore unnecessary.

Hardware Description

RLIN runs on IBM-compatible mainframes. The disk subsystem consists of IBM 3380s in the central site configuration, but it could also be any disk device supported on IBM equipment. Currently, a polled, multipoint terminal network is managed by multiple PDP-11 front-end processors, connected to the mainframe channel with a locally-developed interface device. The terminals are microprocessor-based CRTs, with the MARC/ALA character set in ROM. The link-level protocol used is HDLC. ASCII terminals are also supported via TYMNET or dial-in. Terminals which support vernacular character sets for Chinese, Japanese, and Korean co-reside on the polled network with more standard terminals. Development of non-roman alphabet capability will begin in early 1984.

Underlying Software

RLIN uses a generalized timesharing system, Stanford's ORVYL, as a transaction monitor. The application is completely conversational, because ORVYL retains user memory across terminal transmissions. The ORVYL file system allows allocation of fixed size data blocks to files; the blocks are randomly retrievable by relative block number. Record management functions are supported by the File Services modules of the SPIRES Data Base Management System.

The application, programmed entirely in Pascal, interfaces to the SPIRES record manager through a layered set of Pascal modules, which provide data base level services. Likewise, the application makes use of another set of Pascal services to interface to the terminals without "knowledge" of the physical characteristics of the device.

ORVYL, in turn, runs under IBM's large-scale operating system MVS/XA. Modifications have not been made to MVS itself, but the job entry subsystem has been heavily modified to run with the Stanford text editor which must accompany ORVYL.

Efforts are planned to begin in early 1984 to move the application components of RLIN to a new software environment, which will be based upon a vendor-supported operating system. This system, designed specifically to support high-volume transaction processing, is installable in a few hours on any IBM-compatible system, and is maintained thereafter by the vendor using a teleprocessing link without recourse to systems programming staff.

Installation and Support

RLIN has not been installed outside of the central site at Stanford. The requisite documentation to guide configuration and installation does not exist. No staff support is currently in place to provide systems engineering support, although considerable application expertise exists and would be available for user consulting. User documentation exists for those applications now supported. Library operations specialists are available to perform initial and continuing training of system users.

System Operation

The IBM operating system together with the Stanford software provides the requisite operator functions. A great deal of effort has gone into the automation of production control functions necessary to schedule and run batch production jobs. The new operating system being contemplated would simplify substantially the operational functions now in place. Operating documentation and procedures for this new system will be written in a general way, with any Stanford-specific material clearly separated.

System Administration

The required functions to support privilege definition are present for user identifiers and groups of identifiers, but not for terminals or groups of terminals. Tailoring functions do not exist for defining or modifying terminal display formats. Alterations to tables of code abbreviations and their standard expansions is allowed. Profiling capability exist to support technical processing functions only.

Logging of commands may be turned on or off; if on, a percentage of sessions may be randomly selected for sampling. Response time is captured for each command; resource utilization may be captured as an option. Complete daily and weekly reports are generated; other reports such as session chronologies are available on demand. A software terminal simulator is under development which will allow the placement of arbitrarily large loads on the system for testing and measurement purposes. All of the other performance reporting tools specified in the RFI are supported.

General User Interface

The RLIN primary terminal is a block-mode CRT, which presents a screen with protected element identifiers and unprotected adjacent areas containing data to be modified or containing blanks to be filled in by the user. A unique feature of the terminal is the expandable field, where a movement of the cursor into column 80 signals the terminal to expand the field by one line. The system re-prompts with any data that was pushed from the screen. The MARC/ALA character set is supported, with diacritical marks displayed before the character they modify. Printers may be attached to the terminals, either to print worksheets from screens, or to serve as network printers receiving batch output from the host computer.

ASCII terminals are supported secondarily, in line-by-line mode for searching only. The application is not cognizant of the difference, since it is separated in a well-defined way from the terminal device software. (The same function is supported in other systems by use of a protocol converter.)

RLIN also provides for terminals capable of handling the input of Chinese, Japanese, and Korean characters. These special terminals follow the same protocol conventions as other terminals on the network; the data base software and the terminal software handle the 3-byte character codes in a manner transparent to the application.

The command language is composed of common words and their three-character abbreviations; several commands can be chained together and issued at once. Help facilities do not exist, since the system has been used almost exclusively by trained personnel to date.

Data Base Searching, Retrieval, and Maintenance

All U.S. MARC formats are supported. Care is taken to assure that new formats or changes in existing ones are fully supported in a timely fashion. A library using RLIN as a local system may, within limits, designate the indexes to be maintained by the system. Further tailoring is allowed insofar as specification of fields and subfields contributing to an index and choice of the algorithm to be used in generating the index (e.g. phrase, word, personal name, etc.).

Currently, indexes are updated overnight in the event that a record is added, modified, or deleted. The record itself in its updated form is available for display and further modification on the day of the change if the user specifies the unique key of the record. Additions to the data base software and the record manager would have to be made to support real-time updating of indexes for applications requiring it.

In the central data base, bibliographic records relating to the same edition of a work are "clustered," and RLIN provides transitions from display of the primary record of a cluster to display of a particular record for a particular purpose. A local version of RLIN would deal automatically with clusters with one member.

On-Line Catalog

The current system provides no help or tutorial functions. A command mode, together with automatic prompting of the next most likely command, provides an efficient interface for the trained user. The system is structured to facilitate the addition of tutorial and help facilities should the need become apparent.

Cataloging

RLIN now provides a full-function cataloging system, but without authority control. In terms of the RFI, it does not now provide for global change of identical field contents within records of a search result, and it does not write deleted records to an archive file.

In the first quarter of 1985, an authority control system will be installed to meet the requirements listed in the RFI. In particular, it will support online validation of headings; those maintenance functions necessary to maintain an authority file with logical integrity; and change functions to allow bibliographic records to remain synchronized with the authority file.

A generalized report generation capability now exists to create lists in the format of a page-form catalog.

Acquisitions

The RLIN acquisitions system supports procurement, both unitary and extended. All functions listed in the RFI for procurement are supported. Invoicing and fund accounting are not currently supported, but external design specifications for those functions have been completed and approved.

Serials Control

RLIN does not presently support a serials control capability.

Circulation

RLIN does not presently support a circulation control capability.

Interlibrary Loan

RLIN supports an integrated interlibrary loan system that satisfies the requirements in the RFI.

Costs

It is difficult to be exact about costs in a situation in which some applications have not yet been developed, and in which an operating system change is going to take place. A reasonable expectation, however, would be costs in the same general range as those for BLIS.

Appendix A

RLG Member Activity Profile

Statistics in this appendix represent activity estimates within the RLG member libraries. The numbers given are (small, average, large) taken across the membership.

Circulation	
180K, 690K, 2000K	Annual number of charges
21K, 83K, 200K	Average number of pieces simultaneously in circulation
12K, 53K, 90K	Number of persons holding library privilages
0%, 2.1%, 7.5%	Expected per annum rate of change in circulation activity
Cataloging	
13K, 45K, 110K	Number of titles cataloged annually
5%, 16%, 30%	Percentage of original input
1K, 5K, 12K	Number of added copy/added volume changes to holdings
1K, 9K, 15K	Other modifications to bibliographic descriptions
Acquisitions	
7K, 30K, 107K	Number of items ordered annually
3K, 32K, 116K	Number of items received as direct accessions annually
50, 227, 700	Number of active fund accounts maintained
0%, 1.3%, 4%	Annual expected increase in acquisitions activity

Serials Control

6.5K, 35K, 140K Current number of active serials

43K, 152K, 300K Number of pieces checked in annually

Catalog Searching

Current searching estimates on card catalogs in RLG institutions are somewhat suspect, given the difficulty of gathering hard data. They are reasonable consistent with those few instances where direct measurement has been done.

1K, 8.7K, 20K Number of catalog searches per day

5 Peak/average ratio, hourly rate[4]

Estimates of the annual growth in searching are even more suspect. The following may be used to gain a rough idea of possible growth patterns.[5] It reflects institutional plans for local area networks, and the number of devices that will have access to the on-line catalog.

36, 1054, 4000 Devices installed by September, 1985

41, 2378, 12000 Devices installed by September, 1987

On-line Catalog Size

134K, 487K, 1200K Number of bibliographic descriptions now in machine-readable form

25K, 514K, 2500K Titles that will undergo retrospective conversion over the next eight years

---

4. Taken from a Peat, Marwick, and Mitchell study conducted at New York University. This number represents a single institution, rather than the entire membership.

5. Taken from the responses of eleven RLG institutions to a member survey questionnaire on local area network planning.

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART V

TECHNOLOGY ASSESSMENT

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Table of Contents

Chapter 1 Understanding Hardware Trends 1

1.1 Connectivity 1
1.2 Composition Determines Cost 2
1.3 Limits to Circuit Integration 3
1.4 Cost Implications of Old Technologies 4
1.5 The Learning Curve 5
1.6 Organization of the Assessment 5

Chapter 2 Processors 7

2.1 Comparing Processors by Function 9

2.1.1 Objects and operations 10
2.1.2 Memory model 11
2.1.3 Concurrency and Resource Sharing 12
2.1.4 I/O model 12
2.1.5 Instruction execution speed 13
2.1.6 Parallelism 13

2.2 Cost Trends 16

Chapter 3 Mass Storage 19

3.1 Direct Access Mass Storage 20

3.1.1 Magnetic Disks 23

3.1.1.1 Transfer rate. 23
3.1.1.2 Recording density. 23
3.1.1.3 Track density and seek time. 26
3.1.1.4 Reliability. 27
3.1.1.5 Medium removability. 27
3.1.1.6 Archival Quality of Media. 28
3.1.1.7 Magnetic Disk Price/Performance 28

3.1.2 Digital Optical Disks 32

3.1.2.1 Transfer rate. 33
3.1.2.2 Data capacity. 33
3.1.2.3 Seek time. 34
3.1.2.4 Reliability. 34
3.1.2.5 Medium removability. 34
3.1.2.6 Rewritability. 35
3.1.2.7 Archival qualities of the medium. 35

3.2 Magnetic Tapes 36

Table of Contents

Chapter 4 Printers 39

4.1 Impact Printing 41

4.1.1 Low-speed Serial Printers 41
4.1.2 Line Printers 42
4.1.3 Dot Matrix Impact Printers 42

4.2 Image Transfer Printers 43

4.2.1 Electrophotographic Printers 43
4.2.2 Dielectric Printers 44
4.2.3 Other Transfer Techniques 44

4.3 Treated Paper Printers 45

4.3.1 Thermal Printers 45
4.3.2 Electrosensitive Printers 45
4.3.3 Dielectric Printers 45

4.4 Ink-jet Printers 46

4.4.1 Electrostatic Steering 46
4.4.2 Magnetic Steering 46
4.4.3 Multiple Jets 47
4.4.4 Drop on Demand 47

4.5 Cost Trends 48
4.6 A Note on Storing Digital Images 50

Chapter 5 Information Displays 53

5.1 Cathode Ray Tube Displays 54
5.2 Alternative Display Technologies 56

5.2.1 Flat CRT Displays 57
5.2.2 Plasma Displays 57
5.2.3 Electroluminescent Displays 58
5.2.4 Liquid Crystal Displays 58

5.3 Control Electronics for Graphics Support 59
5.4 Health Impact of the Video Display Terminal 59

5.4.1 Current Knowledge 60
5.4.2 Future Trends 62

Table of Contents

Chapter 6 Communication Hardware and Transmission Channels 65

6.1 Layered Architectures 66
6.2 Long-Distance Transmission Media 68

6.2.1 Telephone Company Services 69

6.2.1.1 The telephone network. 69
6.2.1.2 Digital vs. analog signalling. 69
6.2.1.3 Modems. 70

6.2.2 Value-added Networks 71
6.2.3 Satellite Transmission 73

6.2.3.1 Bandwidth. 73
6.2.3.2 Reliability. 73
6.2.3.3 Operational complexity. 74
6.2.3.4 Cost. 74

6.2.4 Other Providers of Subscriber Connections 74

6.3 Local Communication 75

6.3.1 Network Topologies 75

6.3.1.1 Star networks. 75
6.3.1.2 Bus networks. 76
6.3.1.3 Ring networks. 76
6.3.1.4 More complex topologies. 76

6.3.2 Transmission Media 77

6.3.2.1 Twisted wire pairs. 77
6.3.2.2 Coaxial cable. 77
6.3.2.3 Optical fiber. 77

6.3.3 Baseband vs. Broadband Local Area Networks 78
6.3.4 Access Protocols 78

6.3.4.1 Contention protocols. 79
6.3.4.2 Distributed collision-free protocols. 80

6.3.5 Commercial Local Systems 81

6.3.5.1 Baseband LAN's 81
6.3.5.2 Broadband LAN's 83
6.3.5.3 Digital PABX's 83

6.4 Cost Trends 84

Table of Contents

Chapter 7 Software Engineering 85

7.1 Software Quality 85
7.2 Software Life-Cycle 89

7.2.1 Requirements Definition 90
7.2.2 External Design 90
7.2.3 Internal Design 91
7.2.4 Implementation 92
7.2.5 System Integration and Testing 94

7.3 Development Tools 95

7.3.1 Programming Languages 95
7.3.2 Program Editors 96
7.3.3 Symbolic Debuggers 96
7.3.4 Source Control Utilities 97
7.3.5 Development Support Hardware 97

Chapter 8 Operating Systems 99

8.1 Introduction 99
8.2 Requirements of an Operating System 100

8.2.1 Transaction Management 100
8.2.2 Data Base Management 100
8.2.3 Dictionary Management 101
8.2.4 Distributed Processing 101
8.2.5 Other Requirements 102

8.3 Survey of Current Operating Systems 102

8.3.1 UNIX 103
8.3.2 VMS on the VAX 103
8.3.3 MVS on 370-Compatible Machines 103
8.3.4 VM on 370-Compatible Machines 104

8.4 An Assessment of the Future 104

8.4.1 Research and Development Directions 104
8.4.2 Vendor and User Reluctance 105

Chapter 1

Understanding Hardware Trends

> Connectivity will pace the development of computer technologies in the coming decade. Connectivity is more than just networking: it takes into account the number of devices one can connect, the speed at which they communicate with one another, the distances over which they communicate, and the number of different types of devices.
>
> --Robert Metcalfe, 3Com Corp.

1.1 Connectivity

We can better understand what advances to expect in distributed systems in the eighties by imagining such a system exploded to show its technological components. At the highest level, a distributed computer system consists of many individual computer systems, each connecting to one or more of the others by a communication medium and exchanging information according to agreed upon protocols.

Although the single systems need not be identical, each can be further decomposed into hardware subsystems (processors, storage devices, printers, displays, and input devices). Each hardware subsystem, in turn, depends on underlying hardware technologies (analog electronics, digital electronics, materials science, industrial engineering).

The most elaborate hardware assemblage does nothing useful unless appropriate software controls it. This software can be decomposed into software components (operating system, resource managers, tools, and applications) that depend on software technologies (software engineering, database management, operating system design, automata theory, pattern recognition).

Virtually all computer hardware and software research and development efforts are characterized by their efforts to achieve connectivity. Connectivity has three aspects: parallelism among components, bandwidth between components, and elegant and simple component boundaries. Parallelism means the ability to employ several components of a given type if the power of one such component is insufficient. Bandwidth, a term borrowed from communications, denotes the

amount of information that can flow across a connection in a given time. Component boundaries define the rules by which an element connects with other parts of a system. For a component to be easy to use, its boundaries must be drawn to minimize the need to know in detail how the component works. These three aspects form the common theme for recent and expected technological advances. Obstacles to achieving parallelism, bandwidth, and elegant and simple component boundaries characterize the technological challenges for the rest of the decade.

1.2 Composition Determines Cost

Processors, memory, disk drives, tape drives, printers, terminals, and communication interfaces: these are the pieces of a computer system that we buy separately and that often are packaged separately. To understand how changes in technology affect their function and cost we should understand the pieces that compose any computer subsystem.

Any subsystem includes some or all of the following elements: digital control electronics, volatile information storage, non-volatile storage, electro-mechanical parts, a means to display status and receive human instruction, analog circuitry, power supplies, a cooling mechanism, packaging, and the labor to assemble it. By looking at the percentage of the subsystem made up of each element (for example, the proportion of electromechanical parts) and considering the cost trends for that element, overall cost trends can be discerned. In addition to possible technical advances within a technology, cost-effective replacement technologies and the introduction of automated manufacturing techniques lower cost.

Digital and analog circuitry consists of discrete components (resistors, capacitors, relays, inductances) and integrated components (medium, large or very large scale integrated circuits) each having metal legs soldered to traces on printed circuit boards. The traces take the place of connecting wires seen in older electronic products. On a smaller scale the integrated circuits, tiny slices of silicon containing many digital circuit elements with fine wires connecting the chip to the metal legs of its package, themselves contain complex portions of the digital electronic circuitry. In particular, volatile memory is now entirely implemented using VLSI components, and some technologies for non-volatile storage also appear as integrated circuits.

If all of the electronic circuitry fit on one PC board, that board will have some kind of connector soldered to it, allowing other parts of the assembly to be driven by the electronics on the board. If the electronics require more than one PC board, each board will have an edge that plugs into a special connector mounted on a motherboard or backplane. A board interconnection protocol is implemented in circuitry residing on each board. The most common such protocols are parallel "buses". In such systems several electrical signals, representing data bits, address bits, and control bits, run across the backplane in parallel, allowing each board to recognize requests that it must satisfy and to respond by placing information on the bus (parallel

signals) in turn to satisfy the request. Multiple cabinets (multiple backplanes) in close physical proximity might be connected with ribbon cables carrying the bus signals between them.

In more complex systems, boards may connect to multiple buses. Alternatively, an entire assemblage may connect externally to the other subsystems, as IBM disk and tape drives connect to their control units, and control units connect to the channel cable.

The card cages reside in racks along with other components: power supplies, fans for cooling, disk drives, printer mechanisms, CRT displays, or whatever else makes up the device. The rack is usually enclosed, and it may have some kind of control panel, with buttons and lights for its human operator.

In considering the economics of building such a subsystem there are general rules that apply regardless of the precise function it performs. First, more components means more cost. Not only do the components themselves cost money, but they use more real estate on the PC board, making it larger, requiring more cooling and a larger enclosure. As the size of the circuit exceeds a single PC board, one must add a backplane and the additional logic on each board to perform the bus protocol between them. Although placing components on a PC board is often automated when the board is manufactured in large volumes, and the soldering is almost always automated, the board must be tested, documented, kept in inventory, and serviced.

This is why such a large cost saving is achieved as more and more circuit functions are subsumed on a single chip. Furthermore, connections from chip to chip on a PC board are slower than connections within a single chip. Connections across a bus are slower still, and are further limited by the fact that there is only one data path that all boards must use to interconnect. Overall, price goes down and performance goes up as circuit integration increases.

Another key advantage in using integrated circuits comes from VLSI design automation. This productivity aid shortens product life cycles from five years to two years. Developing software to use new hardware now takes longer than the development of the hardware itself.

1.3 Limits to Circuit Integration

Minute chip size, which acts in our favor by reducing board real estate, is a problem when you try to connect the chip to anything else. Only so many wires can be bonded to a piece of silicon the size of a fingernail before no room is left for the circuitry. External connections to the electromechanical devices the electronics control typically consume more power and need larger currents than those available directly on the chip. So some amount of external circuitry is always needed to carry information between the world outside and the chip.

Although many analog devices have been implemented in silicon, most circuits on chips are digital: large numbers of gates connected in complex ways, so that some number of digital signals going into the circuit are transformed according to the rules of Boolean algebra to yield the output signals. The signals in and out of the chips represant binary digits (0 and 1, true and false) and must be within certain narrow voltage ranges and have very small currents if the chip is to function properly.

In an analog signal information is carried by the exact value of some electrical parameter that varies continuously over time. A digital signal carries information in discrete steps. Analog signals can take on high voltages and large currents. Physical limitations such as heat dissipation and the electrical properties of materials have prevented the implementation of many analog functions as integrated circuits. Two recent developments are beginning to lead to big advances in integration of analog functions. First, semiconductor companies have achieved advances in power-switching silicon components. Second, specialized VLSI microprocessors are being introduced designed to allow digital solutions to problems formerly amenable only to an analog approach. This trend is most prominently reflected in the declining cost of such communication subsystems as modems.

1.4 Cost Implications of Old Technologies

It is often tempting to consider only the purchase price of a computer system. Unfortunately, maintenance, energy, and labor costs often exceed purchase prices many times. Looking at electromechanical devices (motors to spin disk drives and move arms across disks, hammers to strike character slugs and cause printed output to occur, and fans to draw hot air away from the electronic components) you immediately realize that you have stepped out of the high tech world and into the world of well understood technology, where frequent breakthroughs cannot be expected.[1] Not only will costs associated with this older technology not come down, but because it is less amenable to automated manufacturing and more subject to wear and the need for service, its life cost can be expected to grow. So the proportion of a component's cost that comes from electromechanical components is going to get larger, as long as that technology is used.

In these areas we should look for upcoming replacement technologies or for the possibility that a single electromechanical assembly can go farther, as might happen if a single disk drive had a much larger capacity. Displays, keyboards, and printers incorporate many technologies ripe for replacement. These will be considered in detail in their own sections to follow.

---

1. A recent exception is the DC brushless microprocessor-controlled synchronous motor. This is not a new idea, simply one that is made practical by cheap microprocessors and power-switching silicon components.

1.5 The Learning Curve

Electronics industry spokesmen often speak of the "learning curve". This term describes the behavior of product price over its lifetime. When a product is first introduced and has limited competition, its developer can charge high prices. Buyers seek the functional improvement the product offers. This repays product development costs. But as competitors understand the technology and introduce equivalent products, the price goes down and buyers decide between equivalent products on that basis. Thus, mature technologies offer better price performance in the short term.

1.6 Organization of the Assessment

The next five chapters describe various hardware system components in detail; compare various technologies; discuss trends to be expected in their cost; and suggest innovations that may influence the availability of improvements.

The last two chapters deal with software. The first of these describes the progress which has been made in the last decade with respect to the application of engineering principles to software development. The second discusses the current state of the art in operating systems, and how this relates to RLG requirements.

Chapter 2

Processors

This chapter examines not only the processors themselves, but also the memory subsystems and I/O architectures to which the choice of a processor leads.

In the recent past the industry placed processors into four categories: microcomputers, minicomputers, mainframes, and supercomputers. A brief examination of these definitions shows that they are no longer appropriate but exposes new categories by which we may more usefully examine current and upcoming developments.

Microcomputers include those processors implemented on a single computer chip. Until recently they were characterized as executing under a quarter million instructions per second. Their instructions operate on simple data objects such as characters, or at most on integers represented by 16 bits. They don't distinguish between simple instructions that any program executes and more powerful instructions that alter the ability of the processor to communicate with attached devices. Lacking such a distinction, microcomputers cannot safely run several independent programs simultaneously, since the actions of one program may harm the data belonging to another program. A program's instructions and data occupy 64K bytes or less if no additional hardware is provided. Transfers of data to and from the processor's peripherals and memory occur in chunks of 8 bits. Microprocessors are commonly found as single PC board systems, but often also are implemented as several boards connected by one of the popular bus architectures: for example, S-100 or MULTIBUS.[2] Maximum I/O transfer rates range from .25 to .625 million bytes per second (MB/sec). RLG uses Intel 8085 microcomputers to implement its special terminals, the RLG 40 and RLG 90.

Minicomputer processors have in the past been implemented using multiple computer chips, often using bit-slice microprocessors. Minis execute as many as a million instructions per second. The instruction set includes some that are equivalent to the microprocessor's in function, but in addition some minicomputers manipulate more complex data objects, such as 32 bit integers and floating point (real) numbers, as well as encompassing more complex operations on these objects in a single instruction. Minicomputers provide both user mode and supervisory mode instructions. They allow programs to be written that run in a privileged "supervisory" mode, sharing the processor and attached peripherals among several independent "user mode" programs. Each

---

2. TM Intel Corporation

program is limited to 64K bytes at a time, but more than 64K bytes of memory connects to the processor simultaneously. Supervisory instructions permit this memory to be assigned to various user programs safely, so that one program cannot alter the data belonging to another program. Data flows to and from peripherals in chunks as large as 16 bits. Often the minicomputer design is especially fast (compared to its overall processing speed) when responding to the completion of some transfer of information from or to its peripherals. This permits minis to be used in labs to control instruments and in industry to control manufacturing processes. In the past minicomputers were implemented as several PC boards obeying a bus architecture proprietary to the vendor. The most well known of these is the DEC UNIBUS, transferring data at rates up to 1 million bytes/second. RLG uses minicomputers, DEC PDP-11s, to control its communications lines connecting terminals to the central database.

Mainframe processors execute more than one million instructions per second. Their instructions act on larger and more complex objects (32 bit integers and large floating point numbers as well as groups of characters taken together). Programs occupy many millions of bytes of memory, and a wide array of supervisory instructions permit more elaborate memory management and sharing of resources among cooperating programs. Transfers of data between the processor and its peripherals occur in chunks of four or more bytes at a time. Memory attaches to mainframe processors in such a way that several transfers into or out of the memory occur at once, permitting program instructions to be transferred to the processor at the same time that blocks of information flow to and from secondary disk storage. The central RLG database system operates on a mainframe, the IBM 3081K.

Supercomputers execute many million instructions per second and perform some specialized operations (for example, floating point calculations) very much faster than mainframes (greater than 100 million 64-bit floating point operations per second). Furthermore, the supercomputers have specialized architectures allowing additional complex objects (groups of numbers) to be acted on simultaneously, facilitating certain scientific calculations. Several techniques contributing to the speed of the supercomputer are summarized by one word: parallelism. One such technique introduced in the early seventies, instruction prefetch, fetches instructions beyond the current execution point. The prefetched instructions are placed in memory much faster than the computer's main memory. When the processor completes its current instruction, the next instruction is available to it without a slow memory access. Another technique, pipelining, causes several instructions to be partially executed simultaneously, with hardware interlocks delaying operations that depended on the completion of another operation still in progress. A later technique allows an operation such as addition to be performed simultaneously on several data objects. This is referred to as "single instruction, multiple data" or SIMD. Often, supercomputers have very simple but fast I/O paths that send streams of information to fast secondary storage. An auxiliary computer manages this fast storage and sends its contents to ordinary disk stores. RLG has no application for the supercomputer as it is described here.

2.1 Comparing Processors by Function

These distinctions between processor types are no longer very useful. The microcomputer category was really established to recognize the single-chip computer. Functions provided by a microprocessor are limited by the number of circuit elements that can be placed on a single chip. As technology has advanced in the production of very large scale integrated circuits (VLSI) many of the functional distinctions between microprocessors, minis, and mainframes have already disappeared.

We retain the term microprocessor to refer to a machine's implementation as one or a small number of low cost integrated circuits. Looking at them closely, we see several measures of computer function have been used to assign processors to the four categories. It is now more useful to examine these measures individually.

- Objects and operations -- What objects are manipulated by the processor's instruction set and what operations are performed on each object type?

- Instruction execution speed -- What is the processor cycle time? What size chunks pass between the processor and memory when it fetches instructions? Is there a pipeline? A cache?

- Memory model -- How large can a single user program be, and how much data can it address simultaneously? Does the processor distinguish between memory occupied by program instructions and memory occupied by the different data objects a program acts on? How much memory can the supervisory program (operating system) control across all programs? Are there limits on the algorithms that can be used to manage the memory?

- I/O model -- Does information pass to and from the processor's peripherals over a bus, a channel or by some other means? What units of information does the supervisory program see in a single transfer? What processor resources are required to initiate and complete a transfer? Does the model favor devices that transfer data a character at a time, by block or by message?

- Concurrency and resource sharing -- What mechanisms are provided to share the processor, its memory, and its peripherals among independent programs?

- Parallelism -- How many I/O operations proceed simultaneously? Are parallel instruction streams fetched and executed? Do multiple processors execute tasks from a common workload? How much data can be transferred from all peripherals to memory simultaneously?

To examine these categories, we look at current and expected features for one old and two new representative "microprocessors": the Intel 8080, the Intel 8086 family, and the Motorola 68000 family; one old and one more recent "minicomputers": the DEC PDP-11, and the VAX/11 family. Finally, we comment on another family, the experimental computers under study in various laboratories. Although they may only be a gleam in their developer's eye today, they suggest architectures to expect in five years.

2.1.1 Objects and operations

Table 1 shows that the trend for newer machines is to implement a wide variety of data types directly in the hardware, or in a closely integrated coprocessor (associated chip architecturally related to the main processor). This directly reflects the role played by VLSI in reducing the cost of complex electronic circuits. The 8080 (and its somewhat more capable cousins, the 8085 and the Zilog Z80) have been among the workhorses in implementing microcomputer applications in the past five years. They are still far cheaper per chip than their newer counterparts, but for comparison, consider that a Z80 processor chip cost $50 in 1978 and costs $5 today when purchased in small quantities. Today the slower versions of an Intel 8086 or a Motorola 68000 cost $50. Clearly, within the next three years computer control applications that now require a dedicated minicomputer will be performed with one of these equally capable processors.

Another result of the widespread use of smaller processors supporting the whole spectrum of data types is that programs are becoming more portable. In the past one of the most frequent roadblocks to program portability was the fact that the programmer might easily make hidden assumptions about the representation of numbers on his machine. When transported to a different machine, different representations for the numbers cause different results and probably erroneous operation. Difficulties of this sort have long stood in the way of realizing a return on our investment in software. In the next five years, elaborate programming efforts to avoid this sort of pitfall will be unnecessary.

Uniformity in the fundamental operations supported by most computer architectures also simplifies the production of portable compilers and allow compiler writers to focus on code optimization. As better compilers become available, more and more applications will be written in high level languages, with resulting productivity gains.

One way to achieve uniformity is to simplify the hardware's underlying instruction set and rely on high-level language compilers to supply complex operations and objects through procedures using the simple instructions. This approach was proposed and implemented at the University of California at Berkeley and is called the Reduced Instruction Set Computer (RISC). The project demonstrated that such a machine (designed for implementation on a single VLSI chip) in combination with well written compilers can perform better than more sophisticated architectures. However, most hardware manufacturers do not appear willing to follow this approach, because it depends on sophisticated software and the willingness of their customers to use high level languages.

One new company, Pyramid Technology Corporation, has announced a commercial implementation of a RISC. They have relaxed some constraints of the original machine, though, and included complex operations such as multiplication and division that were excluded from the Berkeley machine. The Pyramid system is also implemanted in conventional medium-scale integrated circuits, rather than as a single chip. The system is directed at academic and scientific customers.[3] In summary, widespread commercial adoption of the RISC concept is unlikely in the near future.

2.1.2 Memory model

The 8080 and the PDP-11 have identical user memory models: 64K bytes for program, data, and stack. The PDP-11/45 and 44 are exceptions to this statement. They allow 64K bytes each for program, data, and stack. Memory restrictions arise from the size of each processor's representation of an address. All the other processors shown in Table 1 allow much larger user spaces.

The segmented architecture of the 8086 permits very large amounts of memory to be used for instruction storage, but it is somewhat difficult to use more than 64K each for stack space and other data. Compilers now available for the 8086 family do not, in general, support data areas larger than 64K. The difficulty is not an absolute one; rather, it is difficult to support both an optimized model that represents addresses as 16 bit quantities and the more general large memory model that needs two 16 bit quantities for addresses. Because most compilers for the 8086 are intended to run on the IBM PC, incorporating the slowest processor in the 8086 family, optimization is favored over generality. The 64K limitation does not appreciably decline in importance with availability of the 80286, since the user program still sees discontinuities at 64K boundaries.

This example points up the undesirability of architectural features leading to optimizations so directly manifested in the function available to user programs. An architecture like that of the Motorola 68000 family does not place a sharp discontinuity into the machine model seen by a user program. Thus, while one might optimize for the case in which less memory is available, no sharp boundary occurs at a particular memory level to warp the design of programs. One might expect a better overall quality in compilers written for the 68000, since its instruction set is more regular.

Although the 8086 family will be widely used over the next few years (large amounts of software will be available for it, thanks to the IBM PC), new processors will almost certainly tend toward address space regularity.

---

3. Clifford Barney, "RISC Supermini Has Stellar Performance", Electronics 56 #16, 11 Aug 1983, p.149-150.

2.1.3 Concurrency and Resource Sharing

All modern processors incorporate features to permit resource sharing among concurrent processes. At a minimum, this requires two features: separate supervisory and user modes and memory protection. A trusted program, the supervisor, coordinates sharing of attached peripheral devices such as disks and printers. Memory protection prevents memory allocated to one program from being corrupted by another.

Virtual memory allows a supervisory program to construct translation tables that convert addresses seen by the user program to physical memory addresses. This feature is only now becoming available for the microprocessors and still adds substantial cost to the system. The original use of virtual memory was to share the scarce real memory resource. Today, real memory shows the largest price decline of any computer system component. It is realistic to put more memory on a personal computer than was available in the past to an entire mainframe computer. Virtual memory management is still used to share real memory among large numbers of users, but it also serves a different function. It permits a simpler machine model to be adopted by the user program, since unnatural addressing discontinuities are avoided.

In large systems with many users it has become an article of faith that real memory should be added to the system as needed rather than trying to remove the performance degradations caused by high levels of paging or swapping with sophisticated software or secondary storage devices. This point of view is likely to be adopted in the design of systems of smaller size as well. Therefore, memory management hardware will be added to permit simpler software design, and the software that controls it will embody simple algorithms for paging or swapping.

2.1.4 I/O model

As shown in Table 1, the trend in I/O architecture is toward support of multiple, simultaneous bulk data transfers using software- or firmware-allocated parallel data paths, with virtual memory support provided for direct memory access (DMA) transfers.

In general, details of I/O processing are moving to specialized auxiliary processors. The IBM 308X channel director is a high end example but it is true even within the microprocessor families. The Intel 8089, for example, is a DMA I/O processor designed for operation with the 8086 family. One implication of these developments is that less costly, smaller scale systems are becoming suitable for applications requiring high I/O bandwidth, such as RLG's application for library local processing.

2.1.5 Instruction execution speed

Although machine performance is often characterized in "MIPS" (million instructions per second), it is very difficult to establish this figure exactly. Different instructions require different amounts of time to execute; thus, different instruction mixes produce various MIPS ratings for the same machine. Similarly, the number of instructions executed does not correspond to machine throughput. However, gross comparisons are possible, and members of the 8086 and 68000 processor families perform about as well as members of the VAX family, in some cases as well as the VAX 780, depending on the bench mark, the clock rate at which the 8086 or 68000 operates, and the speed of the memory to which they attach.

2.1.6 Parallelism

Perhaps the strongest trend in processor architecture is the increasing application of parallelism to achieve high throughput. Parallelism takes many forms. The first ones that we examine are extensions of traditional single processor approaches that require little change in the programs that run under them: tightly coupled processors and special purpose coprocessors.

Tightly coupled processors are two or more processors executing with access to a shared memory. In this environment, the supervisory program schedules units of work (for example, separate user programs) for execution on whatever processor is free at the time. This technique works very well if the programs being executed are independent in function. If, however, both programs require access to the same resources, a mechanism to coordinate sharing must be designed. This mechanism, a form of lock applied to the resource by the first program until it finishes and unlocks the resource, also prevents the processor on which the contending program executes from being utilized. This lock contention is a difficult problem to solve in a general way, and it grows more complex as more tightly coupled processors are added. Thus, a hardware configuration that has two five MIPS processors may not truly represent ten MIPS of processing capacity, but, rather, eight or nine after contention has been accounted for.

Coprocessors, architectural features of both the 8086 and the 68000 families, are special auxiliary processors that work in close cooperation with the primary processor. The primary processor's instruction set includes classes of instructions that it recognizes but does not execute. When the 8086 family decodes such an instruction in a program, it simply discards it (does nothing). An attached coprocessor, for example, the 8087 numeric processor, monitors the instruction stream that goes to the 8086 and acts upon the instruction to provide floating point operations. The numeric processor's operation takes much longer to execute than it takes the 8086 to do nothing with the instruction. Therefore, a synchronization instruction is provided that the 8086 program must execute prior to each coprocessor operation. Then, if insufficient time has passed since the last coprocessor operation began, the coprocessor causes the 8086 to wait while it catches up. When 8086 instructions are interspersed with coprocessor instructions, overlapped

execution occurs.

Parallelism has long been achieved between instruction execution and I/O operations, because these latter operations occur in much longer time windows than instruction execution. Coprocessor execution of floating point operations is again an example of an order-of-magnitude difference in time scale, combined with a natural separation of floating point operations from the addressing and control threads of a program's instruction stream. Another example with similar properties is text processing - string comparison, movement, and search operations. Motorola has announced plans to design a string coprocessor, although no date for its release. To achieve more extensive parallelism among adjacent operations within a computer program is still a project for research.

Table 1: Processor characteristics compared.

	PROCESSORS					
	8080 (Z80)	8086 (286)	68000 (020)	PDP-11 (44/70)	VAX-11 (780)	S/370 (XA)
DATA OBJECT TYPES						
Bit field	N	N	N	N	Y	N
Character	Y	Y	Y	Y	Y	Y
String	(Z80)	Y	N	N	Y	Y
8 bit unsigned int	Yc	Y	Y	Y	Y	Y*
8 bit integer	Yc	Y	Y	Y	Y	N
16 bit unsigned int	Yc	Y	Y	Y	Y	Y*
16 bit integer	N	Y	Y	Y	Y	Y
32 bit unsigned int	N	N	Y	N	Y	Y
32 bit integer	N	8087	Y	N	Y	Y
64 bit unsigned int	N	N	N	N	Y*	N
64 bit integer	N	8087	68881	N	Y*	N
Packed decimal	Yc	Y	Y	N	Y	Y
Unpacked decimal	N	Y	N	N	N	Y*
32 bit floating point	N	8087	68881	FPP	Y	Y
64 bit floating point	N	8087	68881	FPP	Y	Y
80 bit floating point	N	8087t	68881	N	N	N
128 bit floating point	N	N	N	N	Y	XA
Instruction set regularity.	LO	LO	HIGH	HIGH	HIGH	HIGH
MEMORY ATTRIBUTES						
Stack in hardware?	Y	Y	Y	Y	Y	N
Code space separable from data?	N	Y	Y	44	Y	Y*
User space segmented?	N	Y	N	N	N	N
Sup. space segmented?	NA	Y	N	N	Y	N
Maximum user space.	64K	1MB(4GB)	4GB	64K	1GB	16MB(2GB)
Maximum physical mem-	64K#	1MB(4GB)	16M(4GB)	.5M(2M)	16M(1GB)	16MB(2GB)
CONCURRENCY						
Supervisor mode?	N	(286)	Y	Y	Y	Y
Memory protection?	N	(286)	Y	Y	Y	Y
Virtual memory?	N	(286)	(68010)+	Y	Y	Y
I/O MODEL						
Commonly used I/O bus	S100	MULTIBUS	VMEBUS	UNIBUS	UNIBUS MASSBUS	IBM Channel
Byte transfers?	Y	Y	Y	Y	Y	N*
DMA transfers?	Y	Y	Y	Y	Y	Y
Protocol limit on aggregate bandwidth		.625MB	?	1MB	1.5MB 5MB	3MB
Number of buses	1	2	2	1	1(1-4	1/256-
Simultaneous transfers						1/chan
I/O auxiliary proc.?	N	Y	Y	Y	Y	Y

* Only indirect or partial support.
c No multiply or divide.
t Temporaries only.
¬ Listed here are architectural restrictions. Packaging of individual systems may place more restrictions on attached memory. The VAX 750 limits physical memory to 2M, the 780 to 12M.
Many systems implement bank switching with added hardware.
+ With the addition of the 68xxx MMU or other hardware.
- Each family includes a member having 8 bit internal path. Members with 32 bit paths have been announced but not yet delivered.

2.2 Cost Trends

Using IBM-compatible processors, we illustrate the profound changes taking place in processor price/performance. By 1985, the equivalent of an IBM 370/158 will be priced at $150,000. In the mid-to-late 1970's, the 158 fetched $1.5 million. [4] Figure 1 arrays IBM list prices (price per MIPS) over the prior five year period, with a superimposed curve representing a 20% per annum decline. Roughly speaking, the decline has been without large discontinuities, and bears out the predicted price in 1985 cited above.

Given that the market demand for computing horsepower and disk storage may expand by 45% per annum [5], and given the competition afforded IBM by the plug-compatible manufacturers (PCM's), we expect the current trend to continue. IBM, and by implication its competitors, will offer a price/performance improvement of 20 - 25%.[6] PCM's prices can be 15% below IBM's because their costs and profitability goals are not as high. [7] Figure 2 portrays price/performance over the next eight years given a minimum of 20% per annum, and a maximum of 30%. The more extreme figure is not out of the question, particularly with respect to smaller processors in the 1-2 MIPS category.

Older equipment that retains utility and remains compatible with new equipment (because of the manufacturer's policy of retrofitting state of the market enhancements on aging products) can have a price/performance edge of 45%.

---

4. Ulric Weil, Information Systems in the 80's, p. 235.

5. Ibid., p. 275.

6. Ibid., p. 275.

7. Ibid., p. 280.

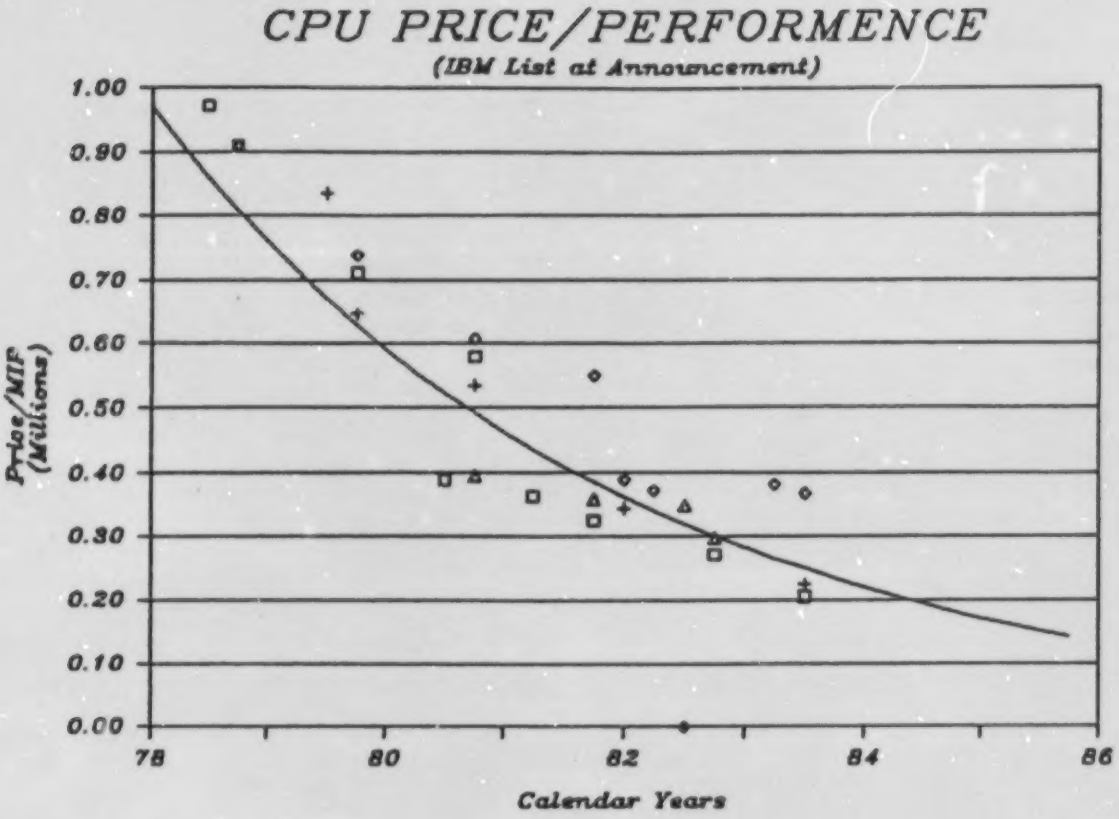
CPU PRICE/PERFORMENCE
(IBM List at Announcement)
Price/MIF (Millions)
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
78
80
82
84
86
Calendar Years

Figure 1

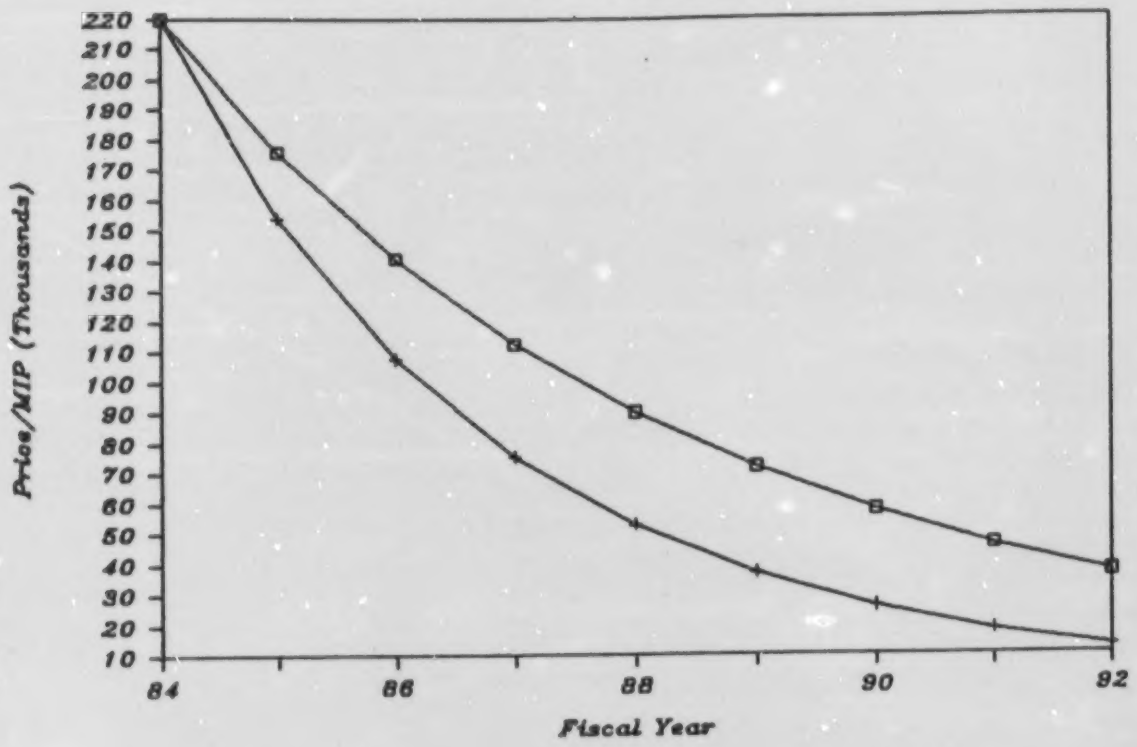
PROJECTED CPU PRICE/PERFORMENCE
Price/MIP (Thousands)
220
210
200
190
180
170
160
150
140
130
120
110
100
90
80
70
60
50
40
30
20
10
84
86
88
90
92
Fiscal Year

Figure 2

Chapter 3

Mass Storage

Advances in storage technology enhance connectivity. That is, much larger amounts of storage are now available, and the stored data is retrievable much more quickly.

In selecting a mass storage device, three attributes of the application distinguish the available technologies: whether direct access to the information is needed; whether the data must be updated after they are first written; and whether the medium containing the data must be removable (for example, to be transported to another site.)

The choice of mass direct access storage in a computer system is dictated primarily by economics. Table 2 shows that now and in the foreseeable future semiconductor storage is at least two to three orders of magnitude more expensive than the most common secondary store, magnetic disk. Yet magnetic disk modules hold two to three orders of magnitude more data than realistic amounts of semiconductor memory.

Table 2: Cost and performance, various direct access storage media over time. (Conservative estimate. Cost, in millicents/bit, is purchase cost, assumed 4x mfg cost. Performance in microseconds/byte). Neither includes combination (e.g.buffer logic) effects.[8]

	1980		1985		1990	
	C	P	C	P	C	P
Fast semiconductor	800	.05	150	.02	40-50	.01
Slow semiconductor	150	.2	40	.1	8-10	.05
Bubble	200	2	50	1,000	5-10	5,000
Magnetic disk	.4	25	.3	20,000	.1-.2	15,000
Optical disk			.1-.05	500,000	.01-.003	500,000

If direct access to the information is not needed, magnetic tape provides mass storage at a still lower cost, as well as enabling data to be removed from the recording device for interchange or archival purposes.

---

8. The Effects of Future Information Processing Technology on the Federal ADP Situation, p. PB82-138181.

Tape uses fall into five categories: processing, interchange, journaling, archiving, and backup. IBM has assessed the tape market in these five areas and concluded that, overall, processing is growing in terms of file size but declining in terms of the number of files processed. Backup and journaling applications have increased; archiving use remains constant; and interchange use has declined with the rise of telecommunications.[9]

Optical disks represent a marriage of the direct access characteristic of magnetic disk and the removability and bulk storage of magnetic tape, but they present an important limitation. Currently available commercial systems do not permit data to be updated once they are written on the medium. This may be an advantage or a disadvantage, depending on the application. Figure 3 shows trends in recording density from 1960 through 1990. Figure 4 shows the relative price-performance of present technologies. Applications such as on-line public access catalogs make heavy demands on mass storage devices in terms of both speed and capacity; therefore, we will continue to monitor trends in this area very carefully.

3.1 Direct Access Mass Storage

Whether a disk is magnetic or optical, direct access to the information stored on it is achieved by a read/write mechanism that, in constant average time, read (or write) anywhere on the surface of a spinning medium coated with a recording material.

Direct access storage devices are categorized by several key features:

- Transfer rate.
- Data capacity.
- Seek time.
- Reliability.
- Medium removability.
- Rewritability.
- Archival qualities of the medium.
- Cost.

---

9. Ed Foster, IBM Corporation, briefing on tape directions, Tucson AZ, May 1983.

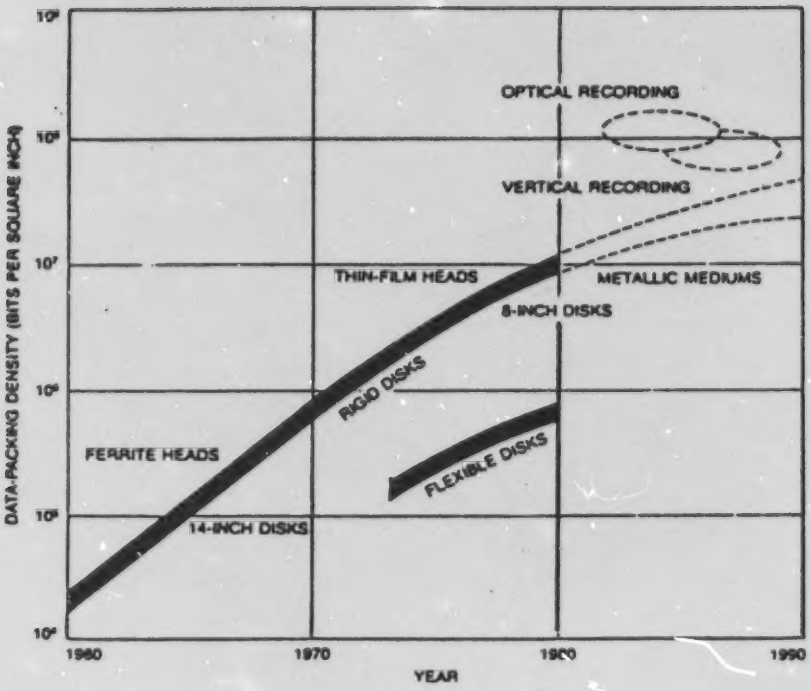

Figure 3

(From "Disk Storage Technology" by Robert M. White. Copyright © 1980 by Scientific American, Inc. All rights reserved.)

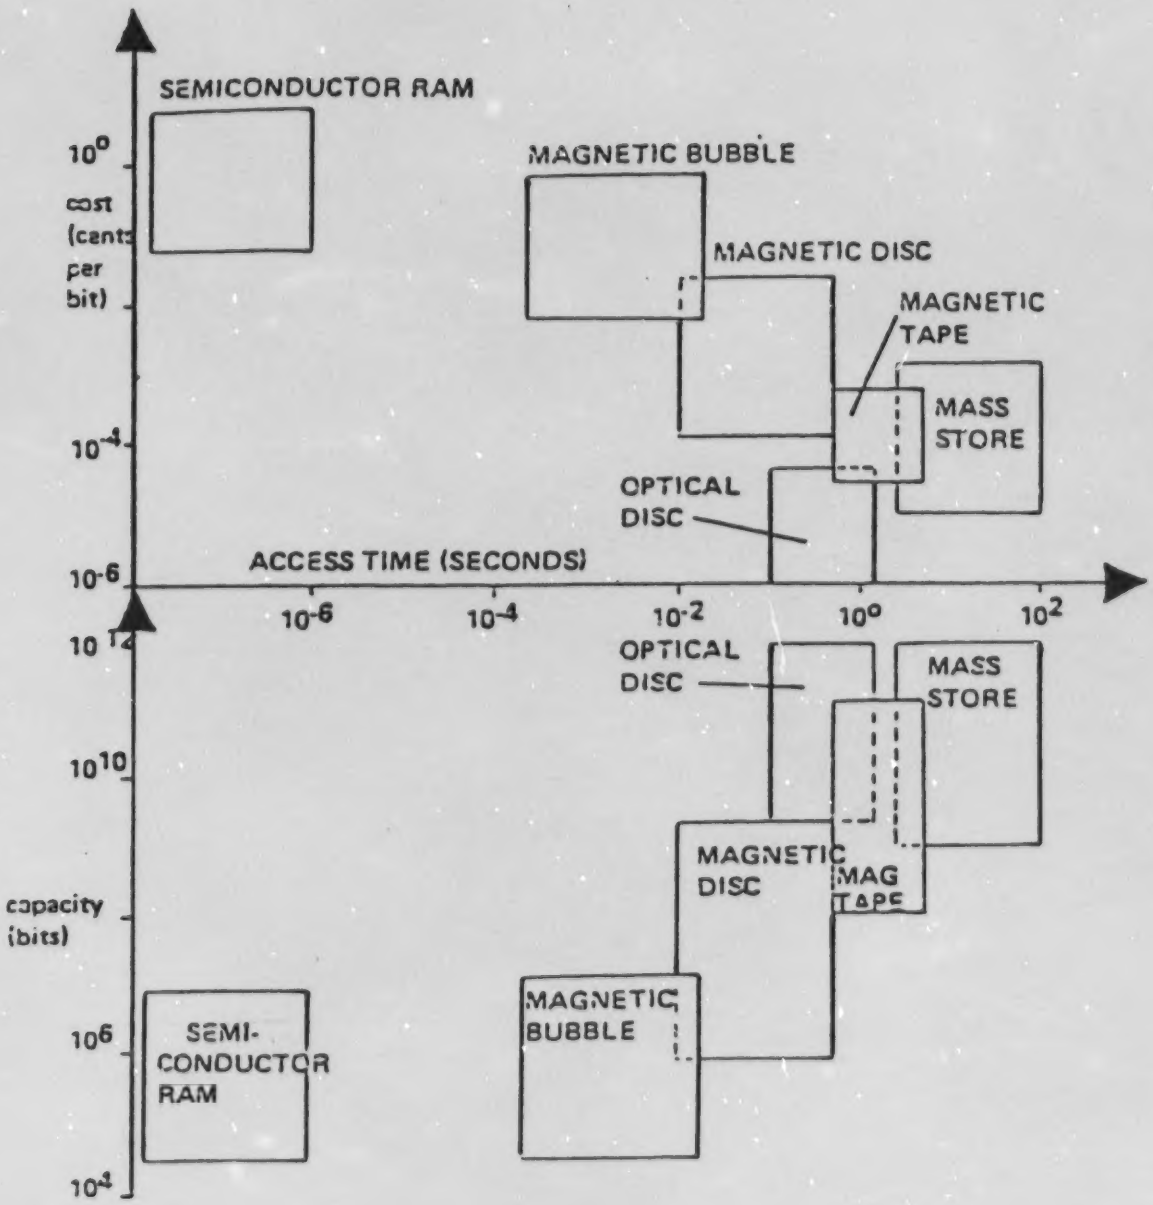
SEMICONDUCTOR RAM
MAGNETIC BUBBLE
MAGNETIC DISC
MAGNETIC TAPE
MASS STORE
OPTICAL DISC
ACCESS TIME (SECONDS)
cost (cents per bit)
capacity (bits)
SEMI-CONDUCTOR RAM
MAGNETIC BUBBLE
MAGNETIC DISC
MAG TAPE
MASS STORE
OPTICAL DISC

Figure 4

Currently, alternatives for direct access storage include 3380-class magnetic disks and digital optical disks for large capacity applications. The chief difference between them is the inability of the surface of an optical disk to be rewritten. For smaller applications designers use "floppy" disks and small Winchester hard disks. A few low capacity applications justify the use of magnetic bubble memory, a form of semiconductor memory, but this technology is only cost-effective in rugged environments (for example, military applications) that preclude the use of disks.

3.1.1 Magnetic Disks

Head and disk composition together determine magnetic properties affecting recording and thus determine capacity, transfer rate, and worst case error behavior for magnetic disks. Notice that this is controlled by recording density and the precision with which the recording and reading mechanism can be positioned over the medium to form concentric tracks.[10]

3.1.1.1 Transfer rate.

Because information is recorded in concentric tracks on the surface of the disk, its rotational speed in combination with linear bit density determines the rate that the drive can transfer data. High rotational speed requires strong disk platters to prevent angular momentum from disintegrating the platter. At the low end of the spectrum of performance and cost, floppy disk platters made of thin sheets of Mylar plastic rotate at 300 revolutions per minute. At the high end, IBM 2305 disk platters made of aluminum rotate at 6000 revolutions per minute. In high performance disk systems, the transfer rates are so high that manufacturers are forced to implement semiconductor memory buffers in the disk controller, because data is transferred from disk faster than it can be sent on a single I/O path to the mainframe.

3.1.1.2 Recording density.

Even if the disk's magnetic coating were able to store information at unlimited density, its ability to do so is limited by design of the read/write head. Ferrite heads consist of a ring-shaped electromagnet (Permalloy surrounded by a wire coil). The ring is positioned as if it were going to roll down the track containing information. Instead, of course, the track rotates below it. The head's electromagnet has to be able to respond quickly to minimize the area it magnetizes on the medium; therefore, one area of research seeks to improve this responsiveness.

The ring has a slot at the bottom through which a magnetic field projects toward the medium. The magnetic field produced at the slot falls off rapidly as the distance from the medium increases. Of course one could let the head rest on the surface (this is how floppy disks work), but the wear this imposes

---

10. Robert M. White, op. cit. p. 138-148.

on the medium makes such a choice impossible at higher rotation speeds. Once the design incorporates a flying head, the technical problems are centered in learning to reduce the flying height. If dirt particles rest on the surface, errors occur reading the data. Heads that fly above the surface operate at such small flying heights that collision with such particles may bounce the head against the surface, damaging it. This event is termed a head crash.

But magnetic media coating the disk are not able to store information at unlimited density. Magnetic disk platters are coated with a slurry of iron oxide. The head, when positioned above the medium, produces a magnetic field, making the oxide particles act like tiny magnets and align themselves in its direction. Visualize the magnetic regions as tiny arrows lined up along the disk track, their heads pointed toward or against the direction of rotation. Data recorded in this orientation (longitudinally) cannot be denser than about 25,000 bits per inch, because the magnetic field of one arrow tends to demagnetize its neighbor.

Domains of opposite direction are not used directly to represent zeroes and ones. Both because of the technique by which the data is read and to provide for error recovery, disk manufacturers use more complex techniques (run-length-limited codes and forward error-correcting codes) to record the data. The advantage in higher bit densities must be traded off against more complex controller electronics to decode the information. Ultimately, decoding becomes so time consuming that the disadvantages outweigh the benefits.

Winchester Technology -- Improvements in disk capacity and transfer rate come from a combination of medium improvements and recording improvements. A major innovation in the design of ferrite recording heads appeared commercially for the first time in 1973 in the IBM 3340 disk drive, and subsequently in the IBM 3350 in 1975.

This enhancement, known as Winchester technology, entails constructing the head in the shape of three parallel rails. This causes the air flowing between them and the disk to act as a bearing, allowing the head to fly at very small distances above the medium. As the distance is reduced between the head and a rapidly rotating medium, other improvements are needed. The metal of the disk itself must be machined to much closer tolerances. The oxide slurry must be applied more evenly to present a smooth surface. Finally, the head and disk must be enclosed in a sealed package, assembled in very clean surroundings, to prevent dirt particles from interfering with the traversal of the head. Despite the solution of all these problems, Winchester heads are still faced with the problem of the durability of the rail material itself. If this material is too thin, it breaks during use.

Winchester technology has only become widely available in low-end products within the last two years, but it has made low cost mass direct access storage with fast access times available for the first time.

Thin-film Heads -- A technology developed for the production of semiconductors has recently been introduced to increase linear bit density by increasing the

responsiveness of the recording head.[11] Thin film heads are constructed by sputtering, evaporating or plating the conductive coil of the electromagnet directly on one end of the air bearing slider. Sputtered on the slider under and over this coil is a Permalloy (nickel and iron) core. This head undergoes much quicker reversals in the direction of the magnetic field. IBM first introduced a disk with thin film heads in 1980. The IBM 3380, using thin-film heads, achieves a linear bit density of 15,000 bits/inch.

Because of eddy currents that cause permeability losses, insulation must be laminated between the layers of Permalloy. This makes the manufacturing process difficult. Though several manufacturers produce thin film heads, only IBM seems to get good yields. These heads are only in use now in the highest performance disk systems and will probably not be seen in low end disk products for one to two years unless IBM enters the low end disk market. This is not unlikely -- rumors have appeared in trade magazines for several months suggesting that IBM may be gearing up manufacturing capacity to build disk drives for the IBM PC.

Thin-film Disks -- Oxide disks exhibit magnetic properties stemming largely from the mechanical properties of the oxide particles that coat them. Past improvements in areal bit density have come from making the particles more dense and more uniform in size. An additional gain, so far demonstrated in magnetic tapes but not used in disks due to durability problems, comes from soaking the oxide particles in a heated solution containing cobalt ions. The particles acquire a layer of cobalt ions, and their magnetic properties improve.

This points up the fact that using a material different than iron-oxide, with better magnetic properties, might improve performance. In fact, a number of companies are making disks from other media to be read with ferrite heads. This combination performs as well as the IBM thin film head-oxide disk combination. Chemically adjusting the cobalt-chromium materials used to coat these alternative disks in a thin-film process improves their magnetic properties. Using this medium, especially in combination with thin-film heads, should almost double the linear bit density to 25,000 per inch.

Disk manufacturers other than IBM have used thin-film media to good advantage to produce high performance disk subsystems to compete with the IBM 3380. Low end products using this technology should also appear in the near future.

Vertical Recording -- Current magnetic disk products employ longitudinal recording, so called because magnetic regions lie with their poles parallel to the disk track. A new technique, vertical recording, has been pioneered in Japan, but is under development by several companies in the U.S. as well, notably IBM. With this technique, magnetic domains point down into or up from the disk's surface. Adjacent domains with the same polarity (arrows pointing in the same direction) tend to reinforce each other. Only at boundaries between groups of like-pointing domains is there a tendency toward

---

11. Jack Taranto, "Thin Film Disks Drive Densities to New Highs", Electronics, April 21, 1982, p. 108-111.

demagnetization. Vertical recording dramatically increases the bit densities possible, with conservative initial linear densities of 20,000 to 40,000 bits/inch. Densities of 100,000 bits/inch have been demonstrated in the lab, and Sony is rumored to have achieved 445,000 bits/inch. The feasible limit, disregarding track positioning limitations, is around 6.5 billion bits.[12] One American company, Vertimag of Minneapolis, is already shipping sample 5 megabyte 5.25" floppy disk drives employing vertical recording.

3.1.1.3 Track density and seek time.

Bit density at which the medium stores data determines disk performance in part, but it is not the only criterion. How accurately the head can be positioned over a particular track also determines capacity. The head positioning mechanism also affects the seek time. Together with rotational delay, this determines how fast a program retrieves randomly chosen elements of data. A magnetic disk's head positioning mechanism has two parts: the device that moves the head between tracks and the method that orients it over the data within a track. Two head movement mechanisms are in use for magnetic disks: stepper motors and voice coil actuators.

Low end disks use the simpler, cheaper, and slower stepper motors. A stepper motor consists of a ring of regularly spaced electromagnets that switch on and off in sequence to rotate a screw that extends or retracts an arm. The disk head is mounted on this arm, so that it moves (seeks) from track to track. Because switching on the electromagnets is relatively slow, the seek time for disks with stepper motors is also slow. More importantly, stepper motors cannot position the head as accurately and yield track densities no better than about 400 tracks/inch.[13]

More expensive, high performance disks use voice coil actuators. These are so called because the mechanism is the same as that which drives a loudspeaker. With this mechanism a circular permanent magnet surrounds a metal rod. A current-carrying wire coils around the rod. When a current is applied to the coil, a force is applied along the axis of the rod, moving it and the disk head across the surface from track to track. The head senses position information while it moves, and this tells it when to stop. Once the head is positioned over the track it must be oriented directly over the data. The head achieves this by reading information and fine-tuning its orientation accordingly. Several techniques are used for storing positioning information: an entire recording medium can be devoted to positioning in very high performance disks, positioning data can be written with application information in the data tracks, or inside and outside tracks can be devoted to

---

12. Tom Manuel, "Vertical Recording Promises New Era", Electronics, April 21, 1982, p.89-90.

13. Larry Sarisky, "Refined Closed-Loop Servo Enhances Low-Cost Disk Drive's Accuracy",Electronics, March 10, 1983, p. 139-142.

positioning information.[14]

Voice-coil actuated disks with closed loop servo-mechanisms achieve the fastest seek times and are able to read and record the highest data density with their accurate positioning. High performance disks such as IBM's 3380 have a closed-loop servo-mechanism, as well as a precision-aligned head. Position information is continuously fed to the head positioning electronics from a medium surface and read head dedicated exclusively to storage and sensing of position information. The most accurate voice-coil positioners yield track densities of around 1200 tracks/inch. Track densities of up to 2,000/inch may be possible. IBM is expected to halve the access time of their 3380 disk in the not too distant future by introducing a model with two actuators for each disk. At present average seek time of a 3380 is <?> milliseconds.

The IBM 3880 Storage Controller with cache buffer demonstrates another way to reduce seek time. The 3880 cache, which can have up to 8 Mbytes of memory, is managed by an intelligent controller that keeps frequently accessed disk blocks in semiconductor memory for fast access. I/O requests issued to disks behind the cache can complete on either channel connecting the 3880, thus allowing twice the raw transfer rate permitted by a single channel.

Use of buffer logic will increase over the next few years, as the decreasing price of digital control electronics permits even low cost peripherals to employ rather sophisticated algorithms. The net effect will be to make the access time of secondary storage approach the speed of semiconductor memory at much less cost.

3.1.1.4 Reliability.

Magnetic disk drives have very good reliability, with a consistent error rate better than 1 in 10**10 bits. Introduction of thin film media has been delayed due to the fact that the medium is somewhat less durable than the oxide medium. However, the most significant error observed with magnetic disks, that of peak shift timing, is markedly improved by use of thin film heads and by the new medium. This should preserve good reliability despite the increased bit densities.[15]

3.1.1.5 Medium removability.

Two problems must be solved for a recording medium to be removable. First, both the drive and the medium must be robust. When the medium is removed, both it and the drive are unavoidably subjected to stresses that a device with a permanently mounted medium never suffers. Second, media must be able to be used on different drives interchangeably. This requirement can be satisfied

---

14. Ibid.

15. Jack Taranto, op. cit.

at several levels. At a minimum, the information must be recorded in such a way that reading the information does not depend on the precise constellation of attributes shown by a particular recording device (for example, positional adjustment of the head).

Before the advent of the Winchester head, almost all disk drives had removable disk packs or cartridges. With the low flying heights introduced by Winchester technology, however, it is imperative to keep the disk medium absolutely clean. As a result, recording mechanisms and disk media are packaged together into a sealed cartridge by workers inside a "clean room".

Today, disk drives with removable packs are still made, but they incorporate old technology, lower performance, and higher cost per bit. Floppy disks are removable, of course, but have very low transfer rates. In the low end of the hard disk market, manufacturers are building Winchester cartridge disks, but many compromises in performance must be made to achieve resiliency. A more serious problem is compatibility. Even if hardware compatibility is attained across a single vendor's drives, there no standards that define the format and electrical characteristics of the data recorded on the medium. Thus, it is not possible in general to read disks on a different vendor's drive, or on the same vendor's drive when used with a different operating system.

Standards are likely to emerge in this area, but not for the next couple of years. In the mean time, removable media are only justifiable in applications using identical hardware and software configurations.

3.1.1.6 Archival Quality of Media.

Magnetically recorded information has a safe lifetime of one-half to two years. Problems of migration and pass-through (in which adjacent magnetic domains affect each other over time) cause deterioration of magnetic media stored for longer periods. Furthermore, removable rigid disks are bulky to store and require a climate-controlled environment. On the whole, rigid magnetic disks are not a good choice for off-line archival of data.

Floppy disks are more compact and demand less rigorous climate control, but their storage capacity is low (currently, a maximum of under 700,000 bytes/disk). Interestingly, some companies pursuing vertical recording technology for floppy disks view their primary application as a backup medium for irremovable rigid disks. As an example, Sony has an experimental dual-sided 3.5" floppy that stores 4 million bytes per side. Compare this with the 91 Mbyte capacity of a 6250 bpi tape reel. When this type of device becomes available at a reasonable price, floppies may be a very appropriate low end choice for short term archival storage.

3.1.1.7 Magnetic Disk Price/Performance

Over the last fourteen years, the price per unit capacity for high-speed magnetic disks has declined by about thirteen percent per annum. During this

same period, deliverable accesses per unit time have improved by an order of magnitude.[16] Figure 5 traces the decline in purchase price per megabyte during this period for IBM disk modules without considering controllers or heads-of-string, whose prices declined similarly.[17] The figure shows a price curve discontinuous at 1975, where Winchester technology became available embodied in the 3340/3350 product line. Price/performance per annum decreases since that time have averaged less than the overall thirteen percent and are closer to eleven percent. It should be noted that the advent of the 3380 with its thin film head did not cause a discontinuity, but rather provided another data point on the current trend line.

Figure 6 portrays a conservative view of the future that takes the "fully loaded" cost per gigabyte (that is, including the controller and head-of-string), taking into account the growth pattern of the RLIN data base and declining it at eleven percent per annum.

The next important announcement from IBM relating to disk technology of interest to RLG is likely to occur in the period from mid-1984 to mid-1985. This is thought to be a double-density 3380 device utilizing thin-film media and more read/write actuators. It is not expected to provide another price/performance discontinuity, however. Later (circa 1988) IBM will probably announce a disk device employing vertical recording. Since order-of-magnitude increases in recording density are possible even without much increase in track density, a price/performance discontinuity may well occur at this point. A second curve is drawn on Figure 6 to portray a factor-of-two improvement that, in our opinion, could well occur.

Although the preceding paragraphs emphasize IBM's disk product line, technically compatible competitive offerings are available from other companies. For equivalent products, these companies prices can be 15 - 20% below IBM's, because their costs are lower and their profitability goals are not as high.[18] IBM, in turn, keeps pressure on these companies by periodically reducing prices.

At the same time, there is continuing strong demand for more disk capacity. This fact was demonstrated when IBM slipped deliveries on the 3380 by nine months in 1981 due to manufacturing difficulties. The resulting shortage in disk devices caused prices for the older 3350 on the third party market to temporarily rise to 150% of IBM list. [19]

In summary, both the demand and the competition is strong; it is neither a buyer's or a seller's market. This, coupled with the discontinuous behavior of price/performance makes predictions in the disk arena somewhat more

---

16. IBM Journal of Research and Development, Volume 25, Number 5, p. 673.

17. Announcement prices provided by the IBM Corporation.

18. Ulric Weil, op. cit., p. 280

19. Computer Economics, February, 1982, p. 4.

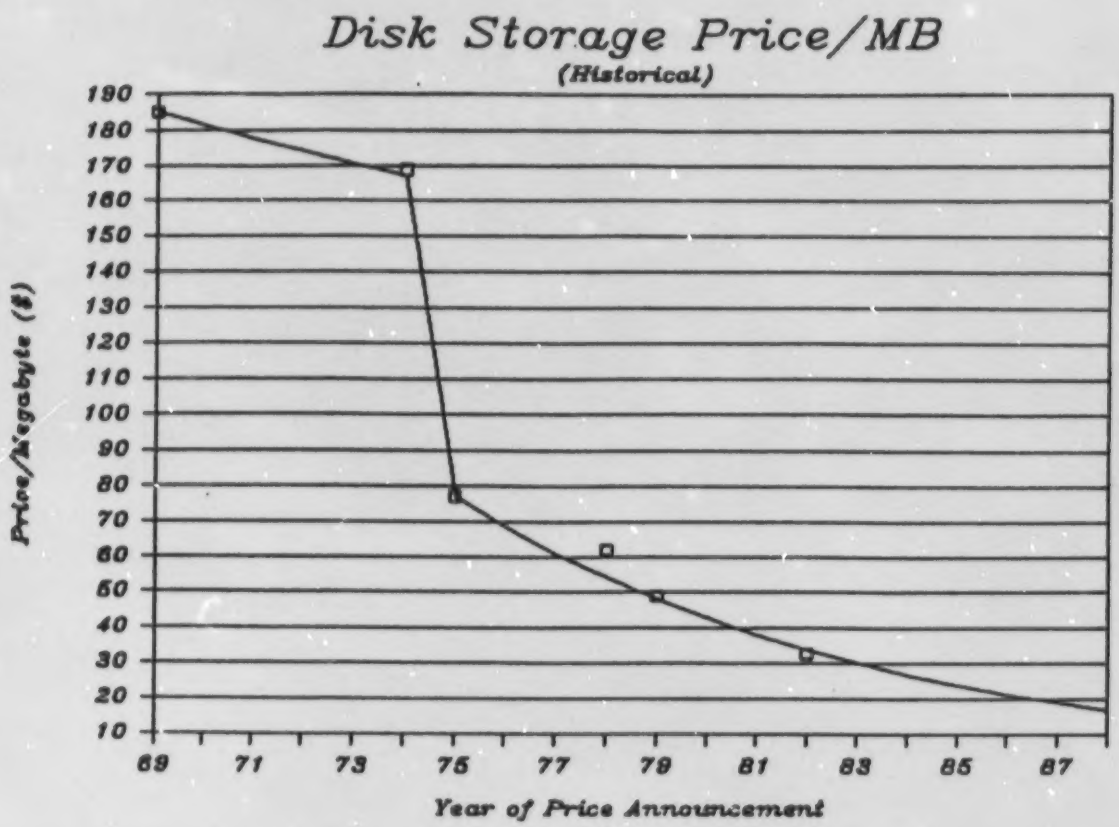
Disk Storage Price/MB
(Historical)
Price/Megabyte ($)
190
180
170
160
150
140
130
120
110
100
90
80
70
60
50
40
30
20
10
69
71
73
75
77
79
81
83
85
87
Year of Price Announcement

Figure 5

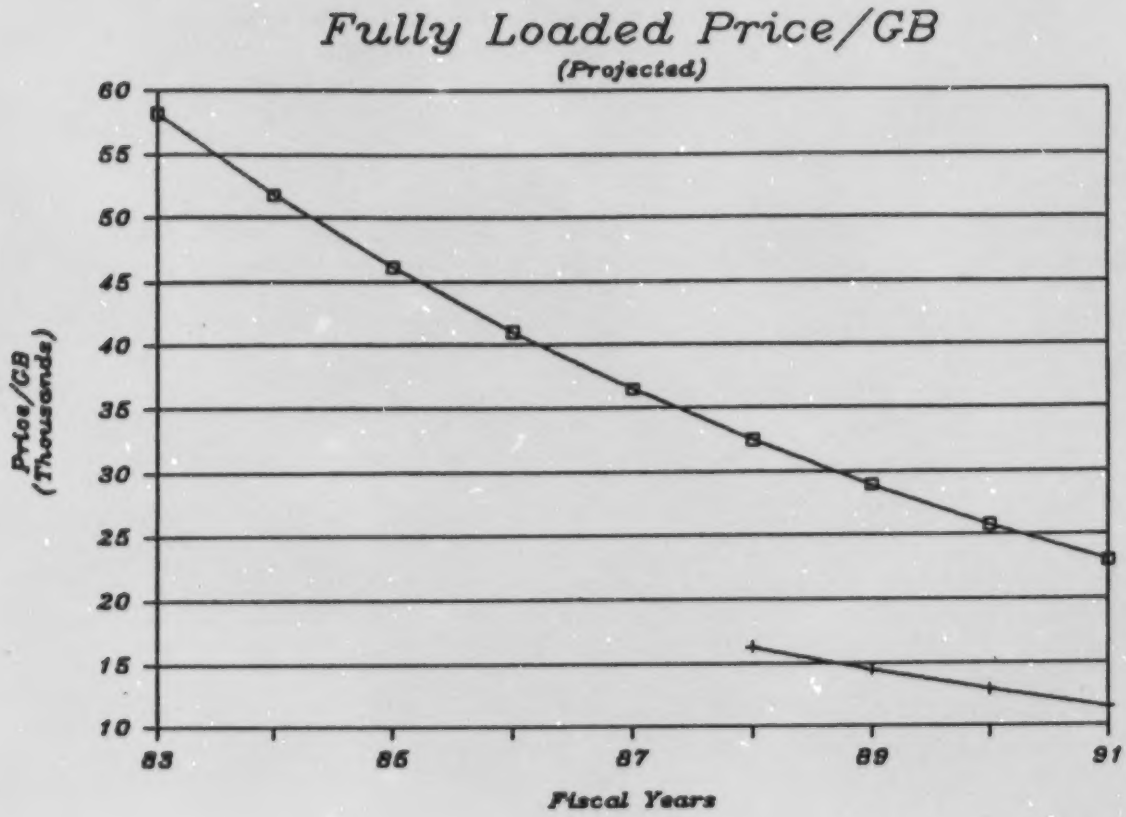
Fully Loaded Price/GB
(Projected)
Price/GB
(Thousands)
60
55
50
45
40
35
30
25
20
15
10
83
85
87
89
91
Fiscal Years

Figure 6

difficult than with processors.

On the low end, economies of scale in the current floppy and Winchester markets should drive prices down somewhat on existing drives, but the largest price-performance improvement should occur with volume availability of high capacity vertically recorded floppies and 5.25" Winchesters in mid- to late 1984. Because the Japanese manufacturers of these products also intend to use them in consumer audio products,[20] expect large economies of scale early in these products' lifetimes.

3.1.2 Digital Optical Disks

Despite the high bit densities attainable with magnetic recording techniques, magnetic disk technology is currently blocked from achieving huge capacity increases due to limits in track density. This is true primarily because reading the information requires that an electromechanical device be placed accurately in close physical proximity to the information to be read.

Another direct access medium, the digital optical disk, achieves linear bit densities equivalent to today's best magnetic disks but is able to take full advantage of the areal density by increasing the number of tracks per inch. It reads information by sensing reflection or transmission of a precisely focussed beam of laser light. Information is recorded by burning a pit in a reflective substance encased in a transparent sandwich layered on a disk surface. When light strikes parts of the surface that have not been pitted, it reflects back. When it strikes a pit, light is absorbed by the dark material of the disk. Some systems may use different materials for coating the surface of the disk or forming its base, but all act on the principle of being able to distinguish between the unrecorded and recorded surface by differing behavior upon exposure to the beam from a laser. Recording densities of hundreds of millions of bits per square inch have been obtained. The good news is that track densities equal to linear bit densities are possible. The bad news is that it is difficult to make the recording process reversible. Although a few erasable systems have been demonstrated, most analysts expect read-only and write-once systems to dominate the market until the late 1980's.[21]

This medium is still in its infancy. However, one large-scale system has been announced, a write-once system made by STC, and several others have appeared to serve office automation applications (see Table 3). Subsequent paragraphs examine some of the problems as well as the promise of this technology.

---

20. Manuel, op. cit.

21. J. Robert Lineback, "Optical Memories Stampede to Market", Electronics, October 20, 1983.

Table 3: RECENT OPTICAL DISK ANNOUNCEMENTS

DATE	COMPANY	PRODUCT
1982	Matsushita	Document filing system.
1982	Toshiba	Three document filing systems.
1982	Hitachi	Document filing system.
Feb 83	3M	Write once medium.
Mar 83	LaserData	Home computer peripheral for read-only disks.
Mar 83	NEC	6300 Intelligent Terminal with optical drive using 3M disk for document storage.
Apr 83	Matsushita	Prototype erasable color-picture file.
Jul 83	3M	Announced year-end availability of erasable magneto-optical medium samples.
Sep 83	DuPont	Announced intent to second-source STC media.
Sep 83	Philips Info. Systems	Megadoc disk and jukebox for 64 12" platters with scanner, display, and printer.
Sep 83	Sony	Announced erasable magneto-optical disks in two years (13cm. disks with 150 MB).
Sep 83	STC	STC 7600 write-once system for IBM mainframes.
Oct 83	Shugart	Optimem 1000 drive (3M medium) for $6000 in quantities of 250.
1984	CDC	Will announce single platter 1GB disk.

3.1.2.1 Transfer rate.

Just as is the case for magnetic disks, the transfer rate for an optical disk is determined by its rotational speed and its linear bit density. The newly announced STC optical disk rotates much slower than their high performance magnetic disks. Because of the high bit density attainable through optical recording, rotation at a higher speed pushes the data transfer rate above the capacity of the rest of the system.[22] Since optical drives achieve high transfer rates with lower rates of rotation, they require less power and induce less motor wear than conventional magnetic disks.

3.1.2.2 Data capacity.

Optical disks have recording densities up to 100 million bytes per square inch. Data capacity is determined by this density, error characteristics of the recording medium, and the ability of the read/write mechanism to position accurately to individual tracks. At present this combination of factors has led to a formatted capacity of four gigabytes on the STC 14" platter, and one gigabyte per side on the 3M 12" platter. Compare this with the 2.5 gigabyte capacity of the IBM 3380, premier among magnetic disk drives.

---

22. John Harper, STC, private communication, April 7.

3.1.2.3 Seek time.

Seek time on optical disks is potentially shorter than for magnetic disks, since movement of the read mechanism occurs differently. Instead, a tiny movement of the mirror used to direct the laser to its proper spot on the disk "seeks" to nearby tracks. STC reports a maximum "fine seek" of 7 milliseconds within a recording band containing six to seven megabytes of information. Seeking across bands is much slower, averaging 50 milliseconds.

3.1.2.4 Reliability.

One of the biggest problems that stood in the way of bringing optical disk to market was the comparative unreliability of the optical recording medium. To correct this problem, data areas on the disk are read back immediately after being recorded. Should an error be detected, the information is rewritten in a new location. Because of the high capacity of the medium, storage devoted to rewritten data and error correcting codes still leaves a data storage area many times larger than conventional magnetic disks. Even after these techniques are used, the error rate is two orders of magnitude higher than magnetic techniques.[23]

One company, at least, does not share the view that reliability problems are solved. The Burroughs Corporation dropped a $20 million development program in April, and company officials give as their reason that "it is impossible to produce a reliable medium for high-performance applications at prices competitive with magnetic storage."[24]

3.1.2.5 Medium removability.

A big advantage of optical disk technology is its ability to be used with removable media. The primary difficulty in removing magnetic media arises from the destructive effect of medium contamination. On optical disks the recording surface is enclosed in a transparent sandwich that holds any dirt particles outside the focal plane of the read mechanism, thus minimizing resultant information loss.[25]

One commercially available system (Megadoc from Philips Information Systems) has already taken advantage of optical disk removability by offering an automated "jukebox" to mount and remove platters. This extends the volume of

---

23. Strategic Incorporated, Impact of Optical Memories (Videodisks), III, May 1981.

24. Lineback, op.cit.

25. Charles M. Goldstein, "Optical Disk Technology and Information", Science, February 1982.

information it supports from the one gigabyte stored on a single side of a single platter to 128 gigabytes on both sides of its 64 platter capacity. Another such system is expected as an accessory to the recently announced STC optical drives.

3.1.2.6 Rewritability.

No commercially available optical disk system is rewritable. Some read disks produced elsewhere, and others write the medium only once. Several manufacturers have prototypes or are working on erasable, rewritable media and drives, however. All systems sold to date force a magneto-optical recording material into a special orientation by applying a magnetic field when a write laser heats the surface. This operation can occur more than once.[26] The read laser operates at a much lower power and thus does not alter the magnetic orientation. Light reflected during the read process discloses different polarizations for zero-bits and one-bits. Vertical recording produces the largest polarization effect.[27]

3M has announced its plans to ship samples of rewritable media at the end of this year. One of the engineers working on the project has predicted that systems using the disk could be available as early as 1985. The new disk has a 1.5 megabyte capacity per side, about the same as the write-once disks 3M already produces.[28]

3.1.2.7 Archival qualities of the medium.

Another early problem in the development of optical recording was the discovery that the most commonly used surface material corroded over time after exposure to air. This problem has been addressed by enclosing the surface in a transparent sandwich. Some manufacturers claim a lifetime of ten to forty years for their medium. Because the disk holds data equivalent to forty conventional magnetic tape reels, it is an attractive alternative archival technology, with better long-term error behavior.

---

26. Karen Berney, "Erasable Optical Disks Could Be Here in 1985", Electronics, July 14, 1983.

27. White, op. cit.

28. Berney, op.cit.

3.2 Magnetic Tapes

In a May 1983 briefing at their plant in Tucson, IBM representatives described characteristics of a new generation tape drive. This device, when available, will set the magnetic tape standard for the rest of the decade. Its release has been rumored for some time.

Criteria for rating magnetic tape performance include:

- Recording speed.
- Medium capacity.
- Data integrity.
- Reliability of the drive.
- Cost.

Recording speed is a product of track density, bit density, and tape speed. Current tape drives record nine tracks across a half-inch medium. Fifty tracks have been achieved in the lab. The next generation tape product will have either 18 or 27 tracks, treated as multiple 9-track groups. Iron oxide media now permit densities of 10,000 to 20,000 bits/inch. Chromium oxide or chromium ion coated oxides soon will allow 30K to 32K bits per linear inch.

Increased track and bit density permit tape speed to be decreased 60 per cent. This change permits a simpler mechanical design, eliminating the vacuum column, minimizing power consumption, and reducing moving parts. In short, a reel-to-reel transport is proposed. Extensive buffering in the control unit will be provided to achieve speed matching for enhanced performance. IBM expects a tape transfer rate of 3 Mbytes/second initially and 6 Mbytes/second by the end of the decade.

Capacity of the new medium is targeted to both a small and large file size. Price ranges, using a current medium price of $15/reel as a comparison, are projected at under $15 for a 70 per cent capacity reel and $19 for a 150 per cent capacity reel of the new medium.

Medium composition and packaging differ from the current standard 10-inch reel. Technical alternatives include choosing the surface and substrate compounds, thickness of the tape, and its surface finish.

Because of data rate and integrity requirements the current standard reel is not appropriate for the next generation of tapes. However, the need to read archive tapes ensures that IBM will continue to support the 3420-type drives for the foreseeable future. IBM is now considering OEM sales of drives and drive components (transports) to permit other vendors to provide

IBM-compatible drives. [29]

Data integrity is influenced by flying height of the head, bit size, medium coating and surface finish. Newer head designs will halve current flying heights of 25 microinches. Thin film technology permits a smaller charge to be used and produces a smaller bit size. New media will be manufactured in super-clean environments, and their use will require changes to operational procedures for cleaning tape heads to minimize debris in the recording/reading area. Intelligent control units will be capable of better error handling at no loss in performance.

Drive reliability increases through the use of electronic motion control, better customer engineering tools and operational changes. Electronic motion control permits the elimination of vacuum columns and capstans. The new product includes a customer engineering maintenance device similar to those currently in use with the 3880 and 3380, to verify tape unit characteristics and pinpoint troubles automatically. Less labor intensive operations procedures offering fewer opportunities for error will be provided through an auto-load feature and a feeder, permitting from five to ten reels to be loaded automatically in sequence. Medium integrity and speed enhancements minimize the frequency and length of reruns.

The new drive and its medium are smaller and cheaper in operation. The drive itself occupies 30 per cent of the floor space and access room of a 3420. The medium package size reduces storage requirements 50 per cent over canister storage.[30] Heat generation is 30 per cent of the 3420, and power consumption only 20 per cent of the 3420. Over the period from 1950 to 1990 tape storage costs will have declined from $10/Mbyte in 1950 to $.60/Mbyte in 1980 to $.01/Mbyte at the end of the decade. Cost of drives expressed in dollars per Kbyte/second have reduced from $113 in 1950 to $.06-$.15 today.

Some evidence for the imminence of announcement of this new tape product is given by the fact that the (relatively new) Tucson facility has extensive thin-film manufacturing facilities and no responsibility for disk manufacture. This point was explicitly made and stressed by the IBM presenters.

---

29. This was expressed as having been suggested in an annual business strategies exercise, not as current policy.

30. Questioning brought out that they were comparing against canisters, not the more space efficient tape seal racks.

Chapter 4

Printers

In the context of distributed processing, RLG needs low-cost, high resolution printers to aid patrons using on-line catalogs, to create copies of materials transmitted electronically in response to inter-library loan, and to distribute information stored electronically, when such storage becomes economical.

Printers are rated on several performance criteria. These include:

- Print speed -- This is measured in characters per second, words per minute or pages per minute. Because some technologies (for example, xerography) produce an entire page at once, comparing their speed to devices that print a line at a time or a character at a time is somewhat difficult. It may be useful to categorize printing as high speed (greater than 100 pages per minute), low speed (10 pages per minute and less), and medium speed (10 to 100 pages per minute).

- Print quality -- Quality is in the eye of the beholder, but it is judged by several concrete criteria. There is a fundamental division between the fully formed character technologies, drawn character technologies, and dot matrix printing. In the first case, characters are transferred to paper from a physical representation of the character: for example, a type slug or a photographic mask. For these printers, quality is measured by the underlying quality of the character representation, the alignment of individually produced characters on the printed page, uniformity of print intensity, and the absence of ghosting or other smudging of the image. Similar criteria apply to printers (for example, many ink jet printers) that draw characters on the page. For dot matrix printers, the appropriate quality criteria include resolution (how many dots are available in the cell that contains a character) and contrast (gray scale), as well as production qualities analogous to the other technologies. Resolution is measured in picture elements per inch. Letter quality output is attained at a resolution of 240X240 elements per

---

31. Where not otherwise noted, information in this chapter was found in Non-impact Markets and Technologies for Distributed Printing, Strategic Incorporated, December 1980.

inch.[31] 600X600 elements per inch is necessary for typography,[32] and the highest quality phototypesetters achieve a resolution of 1800X1800, beyond the limit of human visual discrimination.[33]

- Print versatility -- Printers producing fully formed characters often have only one or a limited number of character representations available. Once multiple fonts are provided, manual intervention may be required to switch between them. An all-points-addressable dot matrix printer can produce any image that can be represented at the dot resolution it provides. Sometimes color images can be produced.

- Forms handling -- Printers vary in their ability to print on untreated paper, paper of varying weight and size, multi-part (carbon) forms, or on both sides of the paper.

- Reliability -- Some printers have adequate reliability at low volumes but do not stand up to continuous use. Printers vary in their need for regularly scheduled maintenance procedures, and in the skill level required to administer such maintenance.

- Environmental requirements -- Noise, dust, and unpleasant odor are three common undesirable environmental effects that may require separate location of printer equipment or provision of ventilation or air filtration.

- Integral intelligence -- Two printers having the same inherent print operations may provide different print quality, versatility, speed or reliability through the presence of integral microprocessor control, buffering or diagnostic programs. Integral control also diminishes the processing required to produce the image by presenting a higher level interface and hiding low level details of printer operation.

- Cost -- This criterion is self-explanatory.

Various technologies are used that provide performance advantages in one or several of these areas.

---

32. Damon Trimble, IBM Corporation, briefing on printer directions, Tucson, AZ, May 1983.

33. Charles Bigelow and Donald Day, "Digital Typography", Scientific American, August 1983.

4.1 Impact Printing

Impact printing is a mechanical process, and consequently will not exhibit much decline in cost over the next few years. Three impact techniques are used to produce printed output: serial, fully formed character printers; line, fully formed character printers; and dot matrix printers. Forms handling is achieved by tractor feed or sheet feed attachments; because of the impact mechanism, these printers directly provide multiple copies through multi-part forms with carbons. Impact printers require only ordinary paper and accommodate various paper sizes and thicknesses. They are usually noisy and thus must be specially enclosed to be used in quiet environments.

4.1.1 Low-speed Serial Printers

These printers hammer a single character image at a time against an inked ribbon onto the paper. The print mechanism advances across the paper until an entire line has been printed, whereupon the paper is advanced to a new line. Three variants of this technique are used. The first, used in the old Teletype, consists of a rotating drum with columns of fully formed character images. The drum is positioned over the ribbon and hammered against the paper to form a character. The second, familiar from the office typewriter, employs a print element comprising images of each character. The print element can be a golf ball shaped device (used in IBM Selectric typewriters) that is tilted and rotated until the selected letter is positioned over the ribbon and paper, then hammered against the ribbon to form a printed impression. "Daisy wheel" printers, introduced by Diablo[34] in the early seventies use a flat, circular print element composed of petal-like spokes, each having a character image at the end. The petals are flexible, and when the correct character is positioned over the ribbon and paper, a hammer pushes the selected petal against the ribbon to form the printed letter. In a variation of the daisy wheel called a thimble, the petals of the print element are folded. Serial impact printers achieve some print versatility by their ability to use different print elements, containing different character sets or fonts. Some position the print mechanism horizontally with a precision of 1/120" and vertically with a precision of 1/48". Then, through the use of a font containing variously positioned dots they can produce graphic output, though at low speed and resolution. They achieve multi-color output through the use of multi-color ribbons. together with a ribbon-positioning control operation. The most sophisticated serial impact printers are capable of forward and reverse motion both horizontally and vertically.

Both the Selectric and the daisy-wheel printer produce very high quality printed output (equivalent to an office typewriter), but both are relatively

---

34. Now a Xerox subsidiary.

slow, with the fastest daisy wheel printers limited to about 60 characters per second. As many typewriter manufacturers have adopted the daisy-wheel technology, cheap versions of these printers have become available. They cost from $500 to $600, run at very low speed (10-12 characters per second), and have few control functions.

In the next few years office quality serial printer price will probably not drop below this level, but more durability, more sophisticated functions (bidirectional printing, graphics mode, multi-colored ribbon support), and somewhat higher speeds (30-60 characters per second) will be available at the low-end price. Expanded use of microprocessor control together with manufacturing economies of scale achieved through sharing common components with office typewriters should bring this about. The technology will soon be superseded by low cost, high resolution impact dot matrix and image transfer printers, however, as is suggested by the fact that Diablo has begun the sale of impact dot matrix and ink-jet printers.

4.1.2 Line Printers

This old technology prints a partial or entire line of printing simultaneously. Three types of line printers have been used since the sixties to produce high volumes of low quality printing: drum printers, band printers, and chain printers. In all three, multiple hammers strike paper against a ribbon to produce fully formed printed characters. In drum printers one complete set of characters rotates (or moves vertically) over each print position in the line. Each position has a hammer that strikes the ribbon and paper as soon as the appropriate character moves over it. Band and chain printers also have one hammer for each print position. A single chain of print elements (often with multiple copies of frequently used characters) moves at high speed above the ribbon. Hammers below the paper strike when an image of the character selected for their position passes above them.

Impact line printers are capable of long-running operation but need frequent cleaning and adjustment. Print bands and chains are sometimes removable, allowing alternate fonts. Chain printers were formerly regarded as very fast. The premier example, the IBM 3211, prints as many as 3000 lines per minute, only medium speed by current standards. For this reason, as well as the cost of producing devices of such mechanical complexity, line printers are largely outdated and are being replaced with dot matrix impact printers at low speeds and page-at-a-time image transfer printers at higher speeds.

4.1.3 Dot Matrix Impact Printers

Dot matrix impact printers print characters a position at a time, in series, across a single line. Several wires of tungsten carbide, each coupled to solenoid magnet, are arranged vertically to form the print head. As the line sweeps across the page, selected magnets are energized to cause the wires to strike a ribbon and create dots on the paper.

This type of printer is the most common low priced printer available today. Popular Computing recently surveyed more than forty printers costing less than

$1000, over thirty of which were of this variety. At higher prices, dot matrix impact line printers are available. For the $1000 to $1500 price, draft output is produced at 240 characters per second (120 to 240 full lines per minute, or two to four full pages per minute) and letter quality resolution of 240X240 dots per inch is produced at one-fourth that speed by repeatedly printing the same line but with a small displacement of the print head.

These printers are rapidly dropping in cost and increasing in resolution. Vertical resolution is limited by the precision with which paper advances. Printers at the highest resolution available within this category provide 960X280 dots per inch (a $3600 printer available now)[35] and 720X240 dots per inch (in early '84 for "under $2000"[36]). A great advantage of the dot matrix approach is its flexibility. Given a mechanism to define the dot pattern corresponding to a letter, support for multiple fonts is assured, and the ability to change fonts under computer control is not much more difficult. Companies offering printers having multiple font capabilities charge for these fonts, as well as for extra features such as proportional spacing or support for alternate alphabets.

4.2 Image Transfer Printers

Each of the technologies in this category forms an image attracting toner particles that are subsequently developed. This image is transferred to paper and ironed or fused to the paper by pressure or heat. The technologies differ in how the image is formed. At present these technologies are more costly than impact printing, but they are also capable of much higher speeds. These printers offer convenience, high quality, and utility at a high price. The difference between copying (where high resolution image transfer technology is available at low cost) and printing arises in the formation of images. As advances in laser technology, VLSI array processors and imaging software occur these differences will diminish and permit image transfer printers to become available for similarly low cost.

4.2.1 Electrophotographic Printers

The most common image transfer printers form their images electrophotographically, as does the Xerox copier. Using this technique, a laser writes the image by discharging areas on a charged belt or drum. Xerox offers three such printers: the 2700, costing $19,000 and printing at twelve pages

---

35. Santec S700 Variflex.

36. Advanced Matrix Technology Office Printer, in "New Products", Electronics, November 3, 1983.

per minute; the 8700, costing $212,000 and printing 70 pages per minute;[37] and the 9700, printing 120 pages per minute. These printers all have a resolution of 300X300. Many proportionally spaced fonts are available for these printers, as is all points addressable graphics. IBM's 3800 Model 3 is also an electrophotographic printer, printing 250 pages per minute at a resolution of 240X240 dots per inch, as compared with 180x144 for the Model 1. When software is available beyond "compatibility mode" support, the Model 3 will have many proportionally spaced fonts available for it, as well as the ability to do medium-range graphics.

This is the most complex printing technology, resulting in reliability problems at high speeds, and few manufacturers have the resources to pursue its core technologies (electrostatics and photoconductivity). Related printing technologies may offer more immediate chances for low cost, fast, high resolution image transfer printers.

4.2.2 Dielectric Printers

A dielectric technique can also construct an image for transfer to plain paper. An ion projection cartridge charges selected areas on an image drum made of dielectric clay. As the drum rotates, magnetized toner particles are attracted to the dot pattern on the drum. Paper and toner are pressed together between the drum and a pressure roller to fix the image. Ions and excess toner particles are then removed. The process is highly reliable. As an example, the Southern Systems Mercurion I printer costs $60,000 and produces letter quality output at 60 page per minute. Southern Systems claims a mean copies between failure figure of 500,000.[38]

This technology has experienced renewed emphasis. The drum coating technique is simpler than in electrophotography and more durable than magnetic. The absence of optical parts results in more design flexibility.

4.2.3 Other Transfer Techniques

Magnetic polarization has been suggested as an alternative way to form the image, since it may allow faster speeds at lower cost and quality. Diablo has recently announced a 300 line per minute thermal transfer printer costing $5000, with a resolution of 200X200 dots per inch.[39]

---

37. "Hardware", Datamation, May 1982.

38. "Hardware", Datamation, December 1982.

39. "Hardware", Datamation, ? 1983.

4.3 Treated Paper Printers

4.3.1 Thermal Printers

These printers, exemplified by the well-known Texas Instruments Silent 700 series, offer reliability at low usages, moderately quiet operation, and portability at low cost. A print head with a matrix of elements, each capable of being heated independently, passes over paper that has been treated with a chemical that darkens on exposure to heat. Three types of print head are used: thick film heads, low cost and heavy duty; thin film heads, where lower thermal inertia yields higher speeds; and silicon mesa, with high element density, limited overall size, and low production cost.

Because thermal printers use treated paper that is unattractive and usually packaged in rolls, their usefulness is limited to applications where draft quality output is adequate and where the higher cost of special paper does not offset the printers' low initial cost. The cost of thermal printers continues to decline, as advances in silicon mesa manufacturing lower the cost of the thermal print head.

4.3.2 Electrosensitive Printers

With this technology, treated paper has a layer removed by electric discharge to reveal a contrasting layer below. The technique evolved from photographic plate making. It is simple and relatively high speed. The paper formerly used was ugly. Quip recently made available paper coated with titanium dioxide that looks better than older papers but is even more expensive ($.10/page as compared to $.02).

In June IBM announced availability in the United States of a 1.5 page per minute electrosensitive printer with a resolution of 600x600 dots per inch. This device prints on roll feed paper with a silver-oxide substrate and is capable of producing half-tones and directly usable press masters.

4.3.3 Dielectric Printers

This process is very similar to dielectric image transfer, except that treated paper is used, making the fixing step of image production unnecessary. Three writing techniques are used. In the first, paper moves under a large fixed linear array of electrodes. In the second, a single stylus scans the page (a much slower method). The third method uses 300 styli on a rotating drum with magnetic toner deposition. The Versatec printer-plotters are notable examples of the first technique. Dielectric printing of specially treated paper is also used by Honeywell PPS and UPI-Unifax (for wire service continuous-tone photographs).

4.4 Ink-jet Printers

Ink-jet printers also produce text from a series of dots, but not all ink-jet printers use dot matrices. These devices paint characters by breaking up ink emitted from microscopic nozzles into tiny droplets and guiding the drops electrically or magnetically to their destination. A gutter intercepts unused drops and recycles the ink. Because the ink drops are very small and can be guided very accurately, these printers produce characters as well-formed as if produced by a 32X32 dot character matrix. These printers are also quiet and capable of machine-directed changes in type font.[40]

4.4.1 Electrostatic Steering

These printers differentially apply a static charge to every other one of a series of continuously generated ink drops and guide their trajectory to the paper by passing the drops between charged plates. The plates' electric field deflects the drops proportional to their electric charge. An example of this type of printer is the IBM 6640. Notice that this approach is analog electrostatic printing, because the position of the droplet stream is guided along a continuously varying path to form each letter.

Problems with this approach arise due to electric and aerodynamic interactions between nearby drops that cause splotching and off-target delivery of the drops. To minimize these effects, only every other drop is charged and elaborate control electronics are employed. This diminishes character print speed and increases printer cost.[41] An alternative magnetic deflection approach avoids some of these problems.

4.4.2 Magnetic Steering

This technique uses special magnetic inks belonging to a class of liquids called ferrofluids, made from a suspension of minute magnetite particles. A stream of ink droplets are deflected through the magnetic field generated by a C-shaped electromagnet. In a series of eight drops moving between the poles of the magnet between the time it is turned on and off, each drop is deflected progressively more acutely, forming a scan line (raster) of droplets. By inserting a smaller, perpendicularly oriented deflection magnet ahead of the raster deflection poles, unwanted drops (white space) are "turned off" by being deflected out of the plane of the raster and into a gutter that

---

40. Larry Kuhn and Robert A. Myers, "Ink-Jet Printing", Scientific American, April 1979.

41. Ibid.

recirculates the drops. Because the magnetic force deflecting the drops is proportional to the amount of material in the drop, larger drops of ink can be used, as might be desired for lower resolution, lower cost printers. (Electrostatic deflection is proportional to the charge applied to the ink drops in ratio to the mass of the drops; thus, it is less appropriate for large drops.)[42]

4.4.3 Multiple Jets

A faster ink-jet printer can be built by employing multiple ink nozzles. By providing a means for each droplet to be independently charged, and by constructing long deflection plates through which all the drops pass, a long raster results. If the charging mechanism acts to produce some or no charge (a binary condition), the deflection plates act to pull charged drops into a discard gutter, allowing undeflected drops to hit the paper. The Mead Corporation has built such a printer, operating at up to 45,000 lines per minute at low resolution.[43] This printer is a multiple nozzle binary electrostatic ink-jet printer.

IBM experimented in constructing multiple nozzles and multiple charging deflectors using silicon fabrication techniques. In these chemicals etch rows of truncated pyramids into a silicon wafer to form the nozzles. The pyramidal shape arises from the crystal structure of the form of silicon used, and silicon was chosen because the technology required for its manipulation is well understood. The charge electrodes are also etched in silicon, with layers of material constituting a semiconductor diode applied, allowing the charge to be applied or withheld. By etching the control circuitry on the charging electrode as well, most of the print element is electronic. In the IBM 6640 ink breaks into droplets through the electro-mechanical action of a piezo-electric crystal. The stream of drops takes on a reflection of the charge of an electrode to which an oscillating voltage has been applied. The periodicity of this charge deforms the fluid's surface, breaking the ink into droplets. Strengthening the oxide layer on the surface of the ink nozzle allows a large enough voltage to be applied to print reliably, leading to a fully electronic printer.[44] Exploitation of this technology could lead to fast, cheap typewriter quality printers.

4.4.4 Drop on Demand

All the techniques described above employ continuous streams of ink droplets, with a gutter to catch unused ink and recycle it. Such a mechanism overcomes problems with starting and stopping the flow of ink and allows fast printing. Another mechanism, drop on demand ink-jet printing, emits ink droplets only

---

42. Ibid.

43. Ibid.

44. Ibid.

when needed by applying pressure impulses, instead of relying on the complicated and expensive ink recycling system. Because this mechanism is simpler (though slower and lower resolution) drop on demand printers are among the lowest priced printers (see Table 4).

4.5 Cost Trends

Over the past few years many low-cost printers have become available, as well as increasingly functional and fast high priced printers. As new products are introduced the trend is toward ever higher speed and resolution for a cost bottoming out at the $300 to $600 level (Table 4). Such a plateau reflects the cost of packaging, powering, and cooling a device otherwise composed of a few very cheap parts. The capabilities of the few cheap parts will increase, but the cost of their packaging will not change. New printers can be expected to be obsolete in two to three years given the current rate of improvement in price performance.

TABLE 4: LOW COST PRINTERS[45]

FULLY FORMED IMPACT	LINESZ	CPS		PRICE
Olivetti Bytewriter O-30	165	12		$ 695
Smith-Corona TP-1	126	12		$ 895
Transtar 130	158	16		$ 895
IMPACT DOT MATRIX	**LINESZ**	**CPS**	**MATRIX+**	**PRICE**
Axiom-Suikosha GP-100	80	30	5X7	$ 266
Data Impact Products DIP-81A	80	100	7X7	$ 369
Commodore 1525P	80	30	5X7	$ 395
Radio Shack DMP 100	80	50	5X7	$ 399
Star Micronics Gemini 10	132	100	9X9	$ 399
Epson MX-80	132	80	9X9	$ 494
Okidata Microline 82A	132	120	9X9	$ 549
Canon A-1200	136	120	16X16	$ 595
Transtar/Seikosha 315	106	50	7X8	$ 599
NEC 8023A	136	100	7X9	$ 645
Star Micronics Gemini 15	233	100	9X9	$ 649
Facit Dataroyal 4510	132	120	9X9	$ 695
Okidata Microline 92	136	160	9X9	$ 699
C.Itoh Prowriter	136	120	NX9	$ 775
Centronics 150 Series	216	150	7X9	$ 795
Commodore 8023P	136	150	6X8	$ 795
Computer Transceiver Execuport	132	80	5X7	$ 795
Integral Data System Microprism	132	110	9X18	$ 799
Radio Shack DMP 200	132	120	9X9	$ 799
Digital Equipment LA-50	132	100	7X9	$ 850
Mannesman Tally MT160I	160	160	7X9	$ 850
Mannesman Tally MT160L	160	160	40X18	$ 990
HI-G 9/80PS	136	150	9X9	$ 995
Micro Peripherals Printmate 150	231	150	7X9	$ 995
Integral Data System Prism	132	200	18X36	$1249
Okidata Microline 93	136	160	9X9	$1249
INK-JET	**LINESZ**	**CPS**	**MATRIX**	**PRICE**
Irwin/Olivetti JP 101	146	80	7X7	$ 499
Canon A-1210	80	40	5X7	$ 795
THERMAL	**LINESZ**	**CPS**	**MATRIX**	**PRICE**
Star Micronics STX-80	80	60	5X9	$ 199
Trendcom 200	80	40	5X7	$ 395
Irwin/Olivetti TH 240	80	320	5X7	$ 850
Radio Shack PT-100	80	50	5X7	$ 995
Texas Instruments 734 KSR	80	30	5X7	$ 995

+ Line size (LINESZ), speed in characters per second (CPS) and character

---

45. David B. Powell, "The Buyer's Guide", Popular Computing, July 1983.

matrix size in dots (MATRIX) are maximums, and are not achieved in combination.

TABLE 5: HIGH PERFORMANCE PRINTERS

IMPACT DOT MATRIX	PPM+	RESOLUTION	PRICE
Advanced Matrix Technology	3	720X240	$ 2,000
Santek S700 Variflex	3	960X288	$ 3,600
IMAGE TRANSFER	PPM	RESOLUTION	PRICE
Diablo Series 200 EPM	5	200X200	$ 5,000
Xerox 2700	12	300X300	$ 18,995
Southern Systems Mercurion I	60		$ 60,000
Xerox 8700	70	300X300	
Xerox 9700	120	300X300	*
IBM 3800 Mod 1	250	180X144	*
IBM 3800 Mod 3	250	240X240	*
TREATED PAPER	PPM+	RESOLUTION	PRICE
Versatec xyz Printer/Plotter			
IBM ...	1.5	600X600	

* Prices for these printers are not directly comparable to those running at lower speed, because they require a much more costly connection to the host computer.

+ Note that for these printers page-per-minute speeds (PPM) are conservatively computed assuming 80 character lines and 60 line pages. Performance would actually be somewhat better for pages with white space.

4.6 A Note on Storing Digital Images

As printers attain resolutions approaching that of typesetting by digital image production, it becomes interesting to consider the direct storage of digital images such as those resulting when a page or photograph is scanned in a digitizer. Note that these images comprise a sequence of bits representing the intensity or gray scale measured at every picture element when the image was scanned. The information is not encoded for machine interpretation as is text that was entered through a keyboard.

Using a maximum resolution of 1200X1200 dots per inch (around the limit of

human visual discrimination)[46], and assuming a compression algorithm of twenty to one, representing an 8.5"X11" page would require 11,200,000 bits. Compare this to a page of encoded text requiring 44,800 bits (assuming 85 eight-bit characters per line and 66 lines per page). The digitized image requires 250 times as much storage, 250 times more time to transmit, or else a 250 times faster communication link. Naturally, if the image contains a picture or some other image difficult to encode any other way more information is also carried. However, to take full advantage of the ability to display printed digital images requires concomitant higher density storage media (as is offered by optical disks), higher bandwidth communication channels (for example, local area network links and optical fiber connections), and high resolution transient displays for previewing images (as provided by bit-mapped raster graphic displays).

---

46. Bigelow and Day, op.cit.

Chapter 5

Information Displays

Information displays provide visual channels to access electronically stored or transmitted data. As the quantity of such information increases, especially in proportion to other channels by which we receive information, it becomes more important that this information channel be of high bandwidth.

Many other characteristics of visual displays are also important. Some of these are technical, others arise from the way human beings see and assimilate information.

Visual displays are compared by the following criteria:

- Viewing ease under different viewing and lighting conditions.
- Gray scale.
- Color.
- Resolution.
- Portability.
- Reliability.
- Cost.

An ideal display is easy to read under different viewing and lighting conditions, is small, rugged, reliable, safe, consumes low voltage and power, is operable over a wide temperature range, and is inexpensive. [47]

Depending on the application, some of these criteria assume more importance than others. Applications can be distinguished as follows:

- Stationary vs. portable -- If the application does not require that the display be carried around, the bulk of the display assumes less importance.

---

47. Frost and Sullivan, The Industrial Display Market in the U.S., July 1980.

- Alphanumeric vs. graphic -- Although alphanumeric displays are implemented by graphically representing the characters to be displayed, great simplifications in the design of displays are possible if all information to be displayed is represented in a relatively small, predefined alphabet.

- Monochrome vs. multi-color -- More problems face the designer of multi-color displays. The application satisfied with monochrome achieves higher resolution at a lower cost.

RLG has many current applications for stationary, alphanumeric, monochrome displays but has an emerging requirement for graphic and multi-color capabilities, with the advent of electronic storage and transmission of text and images.

Many technologies have been tried to satisfy display requirements, and several have had limited success. To date, however, only the cathode ray tube display is fully satisfactory as an information display in the office environment. This chapter first describes CRT technology and trends, and then considers the other technologies, identifying their applicability now and in the light of current research.

5.1 Cathode Ray Tube Displays

CRT displays consist of a sealed glass bottle, an electron gun emitting a focussed stream of electrons, an approximately flat target for the electrons coated with phosphor that glows when the electrons strike, and a yoke mounted around the neck of the bottle that deflects the electrons. The yoke has two deflection systems: one for horizontal and one for vertical deflection. The electron gun operates on an electric current; by altering this current, the beam intensity diminishes or cuts off entirely. Control electronics drive the deflection system to sweep the phosphor from top to bottom in successive horizontal lines. The horizontal lines traced by the beam are called scan lines. By turning the beam on and off very rapidly, individual areas of phosphor glow or remain dark. Thus, the screen can be imagined as a rectangle of tiny dots (pixels). This rectangular dot representation is called a raster.[48]

A pattern drawn on the phosphor by the electron beam soon fades into darkness if it is not refreshed. When a CRT is used as a visual display, semiconductor memory associated with the control electronics contains a binary representation of the raster, and the control electronics refers to this memory to refresh the screen. In the case of graphic displays, each dot on the screen corresponds to one or more bits in the memory, and the writer of an

---

48. William M. Newman and Robert F. Sproull, Principles of Interactive Computer Graphics, McGraw-Hill, 1979.

application directly or indirectly alters all points on the screen. For alphanumeric displays, special integrated circuits have been designed that allow memory to hold encoded representations of each displayed character. When the electron beam reaches a new row of text, the special circuitry selects dot representations for the characters from an auxiliary memory. Thus, the program accessing the CRT can only display the dot patterns stored in the auxiliary memory, but far less information must be transferred to the display when the encoded form of a character is used. This technique is called character generation. Sometimes it is enhanced to allow the auxiliary memory to be modified under program control. In this way, many character fonts or alphabets can be displayed by the same CRT.

Color display is achieved by adding two more electron beams; a metal plate called a shadow mask in front of the phosphor, pierced with a triangular pattern of small holes; and a target comprising three phosphors (red, green, and blue). The three beams are focussed to exactly converge at the holes in the shadow mask so that each phosphor color is struck by only one of the electron beams.[49]

Filters attached to the face of CRTs enhance viewing by reducing glare. Pedestals permit the CRT to be directed in a comfortable position. Brightness and contrast adjustments are often provided.

At the high end, monochrome displays with a 2000X2000 raster are available from Panasonic, and several vendors offer displays of 1280X1024. Manufacturers suggest that 1200X900 is a practical limit for color displays; but no graphics generators are yet available that can take full advantage of these displays, their availability is limited, and their reliability is not good.[50] At the low end, an ordinary monochrome TV monitor has a resolution of 640X480. Low end CRTs represent a mature technology and are highly reliable.

Gray scale is related to resolution, since shading conveys information. If only black and white can be represented, image features small by comparison with dot size exhibit roughness or jaggedness at their edges. By shading edges in varying levels of gray, this problem can be reduced even at low dot resolutions.[51] To represent gradations between black and white, more refresh memory is required for each dot. Representing different colors and hues also requires more refresh memory.

The drastic reduction in semiconductor memory prices, as well as the huge consumer market for television, has led to the widespread use of refreshed CRT displays. The current level of circuit integration allows refresh and character generation logic together with keyboard and serial communication

---

49. Ibid.

50. Jerry Borrell, "High Performance Shapes CRT/Monitor Systems Design", Digital Design, October 1983.

51. Bigelow and Day, op. cit.

line control to be performed with very few chips.[52]

Alphanumeric display terminals are now a commodity, with many sophisticated models available in unit quantities for well under a $1000. Profit margins on such terminals are quite low, and most are manufactured outside the U.S. to reduce labor costs. Thus, further price competition is unlikely. Instead, companies are looking for features that distinguishes their product from others. At first these features can be used to raise prices and increase margins; but, given the competitive marketplace, these features will ultimately appear in the lowest price terminals. As faster microprocessors and larger memories become available, graphic functions are one obvious addition.

CRT technology is not without its problems, however. It is bulky. In monochrome displays it is difficult to achieve more than one to two million resolvable pixels. Color displays are twice as bad, because of the shadow mask and triple beam. At these high resolutions, such a fast refresh rate is required that display flicker becomes a problem.[53] For high-resolution color displays, beam optics, yoke design, and variability between tubes due to aging represent the major unsolved problems. Despite these problems CRT technology will dominate visual displays for the next ten years.[54]

5.2 Alternative Display Technologies

Commercially available alternatives to CRT displays are targeted to market niches with specialized requirements. These requirements include portability, exceptionally large display size (for example, projection TV), ruggedness or the need to use a display in direct sunlight (for example, in avionics applications). Table 6 compares the technologies described in this chapter.

---

52. F.H. Dill, "Future Trends and Mutual Impact of VLSI and Display", 1982 International Display Research Conference, IEEE Society for Information Display, May 1982.

53. Ibid.

54. W.F. Goede, "Technologies for High-Resolution Color Display", op. cit.

Table 6: COMPARISON OF DISPLAY TECHNOLOGIES.[55]

	AC PDP	DC PDP	Hybrid	Flat CRT	EL	LCD
Luminous efficiency	xx		x	xx	xx	n/a
Matrix addressing						xx
Duty cycle				x		xx
Uniformity					x	x
Gray scale			xx		xx	xx
Multi-color	x		x	x	xx	xx
Resolution (40lpi)				xx	xx	xx
Life (color matls.)	x		x	x	x	x
Struct. complexity	x		xx	xx	xx	x
Electr. complexity	xx		xx		x	x
Cost	xx		xx	xx	xx	xx

X = Current problem. XX = Short and long term problem. N/A = Not applicable.

5.2.1 Flat CRT Displays

Flat panel CRT displays employ multiple beams and deflectors. They were originally developed by Northrup and TI for defense applications, but TI has since tabled the product as not offering a large enough market.

5.2.2 Plasma Displays

Plasma displays are constructed of two opposing sheets of dielectric (non-conducting) material covering sheets of glass to which many thin electrodes attach. The opposing sheets are separated by a few thousandths of an inch, and the space between them contains very rarified neon gas. When a voltage is applied to the electrodes, the gas behaves as if it were composed of many tiny independent cells. Each cell can be made to glow or darken independently. In A-C plasma displays a high-frequency alternating voltage applied across the cell maintains the glow. Thus plasma displays have their own inherent memories.[56] Other plasma displays use a D-C technique to maintain the gas discharge. NHK recently demonstrated the first practical plasma color TV display by making a phosphor on one of the glass surfaces glow from an A-C plasma discharge.[57] Similarly Siemans AG recently announced a

---

55. Ibid.

56. Newman and Sproull, op. cit.

57. Charles L. Cohen, "Japan Shows Off Its Flat-Panel Work", Electronics, 6 Oct 1983.

data terminal using cathodoluminescence.[58]

Advantages of the plasma display are its flicker-free image and low bulk. In many ways it is a mature technology. (Burroughs has made D-C plasma devices -- Nixie tubes -- for years, and about forty companies, including IBM, offer products). Its disadvantages are cost and poor luminous efficiency and brightness. Ongoing work in Japan is likely to lead to improvements in these displays for use in large screen, relatively low resolution applications.[59]

5.2.3 Electroluminescent Displays

These displays are similar to plasma displays, except that the material sandwiched between the dielectrics is a powder or thin film electroluminescent material. The most well known electroluminescent display is that made by Sharp and used in the Grid Compass portable computer. The Sharp display sells for over $2,000 in quantity, and contains 512X256 pixels.

Both plasma and electroluminescent displays profit less from advances in VLSI than do other technologies. VLSI does not support high voltages, necessary to switch the excitation electrodes in both technologies. Sharp has made drivers capable of operating in the 200 volt range, but they are not at a very high level of integration. These displays share with other matrix techniques the problem that as integration increases, the ability to add pins to the chip does not. Thus, building a matrix display requires more chips than an equivalent CRT.[60] There is about a ten to one price ratio between flat-panel displays and CRTs.[61] They are not capable of full color display. Their use is probably limited to portable applications and other instances where the size of the display is important, such as in factories.

5.2.4 Liquid Crystal Displays

Liquid crystals are technically composed of "mesophase" materials, occupying a state between solid and liquid matter. These materials are packaged behind a polarizing material. When a voltage is applied across the LCD element, it, too, polarizes light. Because the orientation of the two polarizing sheets is opposite, the activated elements absorb light (they look dark).

Liquid crystal displays show the most promise as an alternative to low resolution alphanumeric displays. Their low power and ability to directly interface to drivers make them amenable to silicon fabrication techniques.

---

58. "Technology", Electro-Optics, July 1983.

59. Goede, op.cit..

60. Dill, op. cit.

61. Frank Catalano, "Exploring Flat-Panel Technology", Mini-Micro Systems, Dec 1981.

Sharp has announced volume production of an 80X16 character display (with 6X8 dots per character cell), 18 mm deep and consuming very little power. A full page display is to follow.[62] Several small liquid crystal displays have been packaged with battery powered portable computers, including the Radio Shack TRS-80 Model 100 (with a 40X8 character display), the Gavilan portable (66X8), and the Convergent Technologies Workslate (?X?). However, the cost of LCD is still much higher than that of CRT displays and the resolution much lower.

5.3 Control Electronics for Graphics Support

Graphic display demands control electronics capable of high data rates and processing power. To alter a screen containing a million pixels in 100 milliseconds requires the computation or transfer of 1.25 million bytes/second. This requires a relatively high performance processor and I/O bus. If the image is sent over a communication channel or stored on disk, these, too, must be capable of high performance. To fully support interactive graphic display requires a dedicated processor running at a high clock rate, supported by fast semiconductor memory. Today, high resolution graphic workstations cost in the range of $10,000. This price will decline as microprocessors of the 3 megaHertz Motorola 68000 variety begin to supplant less capable processors in the lower priced personal computers.

5.4 Health Impact of the Video Display Terminal

A thorough treatment of social issues raised by technological change is beyond the scope of this report. However, one such issue has an immediate, practical effecton the policies of RLG and its members: the real and perceived health impact of video display terminals. This section briefly outlines ergonomic[63] concerns about use of video displays and reports on some research findings in the context of current knowledge and future trends.

---

62. "International Newsletter", Electronics, 21 April 83.

63. Ergonomics is defined as "the science of adapting the work environment to the characteristics of the workers." (Fried, 1982) Within the past few years, ergonomics has gained increasing prominence in the field of computers, first in Europe, and now in the United States.

5.4.1 Current Knowledge

Workers who use display terminals daily complain of visual fatigue, back and other body aches, stress, depression and monotony. With consideration for the special qualities of display terminals and the manner in which they may change job content, these problems can be eliminated or at least alleviated.

Workplace design encompasses lighting, air conditioning, furniture, and general physical surroundings. Display terminals put greater strain on the eyes than print. Characters are sizably larger than normal print and tend to be fuzzier. Often light characters are displayed on a dark background, unlike print. The operator's eyes must adjust each time they move from the display to hard copy. A faulty CRT image flickers, adding to eye strain.

Contrast should be adjusted by the individual operator to a comfortable level, otherwise it causes eye fatigue. One of the greatest contributors to visual fatigue is glare. Bright lights, either from a natural or artificial source, reflect off the screen impairing legibility. The operator may be forced to assume an unnatural posture in order to avoid the glare, and this leads to physical aches, as well. Both Cakir [64] and Sauter [65] have found that terminal operators who wear glasses, particularly monofocals, have a higher incidence of visual and postural problems.

These problems are alleviated by a lower level of lighting, covering windows with curtains (rather than blinds that let in uneven amounts of light), placing a filter over the screen, purchasing a non-glare screen, or using a shield to keep out the light. There are drawbacks to the last three items: filters tend to reduce the clarity of the characters; non-glare screens can be expensive; and shields tend to be clumsy and lead to uncomfortable posture. The operator may be able to move the display to avoid the glare, but the first two remedies -- low level lighting and curtains -- should be the minimum requirement.

As in any office setting, attention should be paid to temperature control. Terminal operators assume a more rigid posture for a longer period of time than traditional clerical workers, and are more sensitive to fluctuations in temperature.

Arndt notes that there are three points of interface between the operator and the workstation: the hand and keyboard; the body and chair; and the eyes and

---

64. Cakir, A., D.J. Hart and T. F. M. Stewart. Visual Display Terminals: A Manual Covering Ergonomics, Workplace Design, Health and Safety, Task Organization. Wiley, New York, 1980.

65. Sauter, Steven L. et al. "Job and Health Implications of VDT Use: Initial Results of the Wisconsin-NIOSH Study." Communications of the ACM. v.26, no. 4. April 1983.

screen [66].

Studies have been done and recommendations made as to what kind of furniture is suitable for a workstation. Contrary to popular opinion, the preferred sitting position finds the operator leaning back slightly with legs extended. Whether the chair should be full-bodied or armless with a high backrest depends upon the tasks the operator performs. But, as Arndt points out[*], the high backrest gives greater support than the small one that is standard on secretarial chairs. The height of the chair and the degree of inclination should be adjustable to fit the individual operator. One study gives the average seat height as 47 cm. [67]

Worktable height should be adjustable to the task and the user. A user of 5'3" should not be expected to use the same workstation as someone of 6'1". The worktable should be adjustable in height, so that the home row of the keyboard is at a comfortable level. Table height ranging from 63 to 74 cm. encompasses all currently marketed keyboards [68]. The German DIN Standard No. 4549 (1981) recommends that desks for workstations be 72 cm. high, or lower if the keyboard is higher than 3 cm. [69]. Worktables must be able to support heavy display terminals, so stability, as well as adjustability, must be considered.

An upright support for source documents, and space to rest forearms and hands both help alleviate physical fatigue. Another feature of the workstation that should be considered in order to lessen visual fatigue is variety of colors and brightness. Colorful posters serve to momentarily rest the eyes from the monotony of terminal displays.

Hardware design should be as flexible as possible, capable of adapting to the physical limitations of each operator. Keyboards should be movable rather than sunk. Muscle tension in the arms, back, and shoulders increases with the distance the operator must reach. One study proposes that "the keyboard configuration may be an important determinant of musculo-skeletal disturbances in display terminal use, with increased stress when adjustability is not possible as with keyboards affixed to the chassis." [70] Keyboard profile should be low and relatively flat. They should be finished in a non-glare matte and have room to rest the hands.

---

66. Arndt, Robert. "Working Posture and Musculoskeletal Problems of Video Display Terminal Operators - Review and Reappraisal." American Industrial Hygiene Association Journal (44), June 83.

67. Grandjean, E. and Vigliani, E. (eds.). Ergonomic Aspects of Visual Display Terminals. 1983.

68. Ibid.

69. Ibid.

70. Starr, Steven J., Claudia R. Thompson, and Steven J. Shute. "Effects of Video Displays," Human Factors, 24 (6), 1982.

Inclination of the display itself should be adjustable by the operator to accommodate individual viewing habits or to avoid glare or reflections. Non-glare screens are recommended. Green or amber displays are easier on the eyes than displays of other colors.

The impact of ergonomics is nowhere more evident than in software design, where great emphasis is placed on being "user-friendly". Before purchasing or designing software it is imperative to understand its user. If the user is untutored, a system-directed approach such as menus is preferable to user-directed commands. To increase ease of use, menus should be fairly short with multiple levels rather than one long list of options. In this way, the user is lead to the appropriate choice in a painless manner.

Job content affects performance and worker satisfaction and so is included in ergonomic studies. Those who use display terminals do so on a number of levels. There are the professionals, creating and composing on the terminal. They may use the terminal a great deal during an average work day, but are not tied to it. Because they are higher level professionals, they shape their work day to a certain degree, taking breaks when they deem it necessary or desirable.

Then there are those who converse with the terminal, frequently or infrequently every day. They have a certain amount of control over their work, which usually consists of a variety of tasks. They are not physically tied to the terminal.

Data entry clerks, on the other hand, have one-way communication with the terminal. Their job is solely to key in data from source documents. They are bound to their terminals and a great deal of emphasis is placed on speed. The better and faster they do their jobs, the more rigid their bodies are and the more restricted their lines of vision. The prime characteristic of their job is monotony and, naturally, this leads to a lack of job satisfaction.

Breaks that were required when doing other tasks allowed for social interaction and the building of a sense of community. This is lacking in an automated office because workers don't have to move about to complete their work. The freedom of traditional office workers to decide when they do each task and to vary their schedule is lost with the transition to efficient workstations. Sauter avers that both job control and social support, two important factors in job satisfaction, are adversely affected in display terminal work.

5.4.2 Future Trends

The United States government has not regulated display terminals as they affect health and safety. Some states have regulated emitted radiation, although study after study show that this is not a problem. Regulations run counter to the need for flexibility and adjustability of display terminals and their workstations. They form obstacles to designing workstations to fit individual tasks and physical limitations.

At present, standards for display terminals are just being developed in the United States. There is a growing recognition that standards for traditional office design often do not apply to display terminals. The European Computer Manufacturers Association has recently recommended to the Common Market that certain standards be adopted. The ECMA recommendations will most likely be adopted in the United States and Canada and so are included as an appendix.

Unions continue to push for established breaks for display terminal workers, but this, in fact, does not address the essential problem of job description and control. Unless managers design jobs to include a variety of tasks, and thereby allow workers some control over what they are doing, fatigue and job dissatisfaction will remain problem areas in automated offices. Another union demand is that pregnant women have the option not to use a display terminal. This was prompted by a series of miscarriages by terminal operators in Toronto. Although studies deny any connection between terminal use and miscarriages, the media and unions continue to bring this concern to the public's attention.

Perceived problems are just as real as established ones to display operators and their superiors. If employees feel that management is ignoring health and safety issues, it damages productivity and a pleasant working atmosphere.

The economics of ergonomics is naturally of concern to everyone. Automation leads to increased productivity, but lack of attention to the human side of automation hurts productivity and the office as a whole. Investment in costly computer hardware is lost if the operator is forced to sit in an uncomfortable chair, or at a table that is too low to accommodate the legs. Operators have a good idea what they are comfortable with and in what kind of a situation they work best.

If managers take the time to speak with the operators, either through surveys or interviews; examine the office environment; and analyze the jobs being done, they prove that their employees are considered valuable assets, and they will also make the wisest investment of capital in designing workstations.

This investment does not necessitate complete redecoration of the office. It can mean buying adjustable chairs and tables or putting up curtains to cut screen glare. It also means that computers are not evaluated in isolation, but as part of the office and an extension of the employee.

Chapter 6

Communication Hardware and Transmission Channels

Data communication occurs over distances as great as around the world and as small as within the cabinet of a computer system. In this chapter we examine long-distance and local communication alternatives. The distinction between them is made by considering whether the underlying transmission facilities must be rented or leased from a common carrier, or whether the physical plant can reasonably be purchased. In general the second option is only practical for networks spanning buildings or campuses, while communication across a city or across the country must use carrier facilities. Irrespective of the distance over which information must travel, hardware choices can be compared based on several common characteristics:

- Bandwidth -- This describes the amount of information that flows across the communication channel in a unit of time.

- Reliability -- Communication channels are subject to various kinds of errors, from brief interruptions that obfuscate a single transmission to more lengthy problems that terminate a connection.

- Transmission delay -- When information travels over thousands of miles (or tens of thousands in satellite transmissions), a significant time may elapse before it arrives. This delay may constrain the functions provided remotely and imposes added complexity on the problem of achieving reliable transmission.

- Operational complexity -- The difficulty of achieving high bandwidth and reliability often leads to the use of complex equipment, requiring close supervision by highly trained operators. Sometimes it may be more cost effective to choose less optimal technical solutions to achieve greater operational simplicity.

- Cost -- As usual, how well a system achieves the other attributes is gauged by considering them in light of the system's cost. Much of the complexity of communication technology arises from efforts to overcome error while achieving high bandwidth at low cost.

The design of a communication system in the light of these performance criteria usually boils down to selecting a "shape" for communication connections (network topology) and choosing rules governing the orderly and reliable flow of information across the connections (communication protocols).

The growing complexity of communication hardware and the increasing use of the technology leads workers in the field to conclude that only a careful definition of protocols permits heterogeneous systems to exchange information reliably. The effort to define these protocols has been an international effort, and occasionally the result reflects anomalies arising from the committee approach to design. Nonetheless, standards now permit systems with well chosen communication hardware and software to link to other systems. That is, a pipeline can now be built between two communicating programs, using standard transport protocols, by which messages inserted at one end can be read reliably at the other. This can be achieved although neither program knows anything about the network through which the data flows.

6.1 Layered Architectures

To structure design of the protocols enabling this end-to-end communication to occur, national and international standards bodies have adopted an architectural model assigning each distributed system function to one of seven architectural layers (see Figure 7).

Some definitions are needed when reading about layered protocols: hosts are the two distributed systems engaging in the communication. Entities refer to the hardware or software components implementing the protocols forming a particular layer. An entity offers a service interface to the layer above and communicates with a peer entity occupying the same position in the layered hierarchy.

In this chapter three layers are of interest: physical, link, and network. These layers, forming the communication subnet, enable the next higher transport layer to offer an end-to-end connection. Unlike higher layers of the model, entities implementing layers in the communication subnet communicate with peers in physically connected parts of the network, rather than with their peers at the destination of the information. The transport layer is the first to engage in end-to-end communication, and thus demarcates communication functions from host functions.

The transport layer uses a service interface offered by the network layer. The network layer provides two main services: addressing (for example, what number to dial to reach the peer host) and routing (which path should the message follow if the network topology offers a choice). If the topology of the network is simple, network layer services may be combined with the next lower layer. Sometimes the network layer is divided into two sublayers. One sublayer performs routing and addressing functions within a single subnetwork, another handles internetworking. This view is useful in organizing complex networks consisting of, for example, disparate local area networks and a packet switched long distance network.

The data-link layer delivers units of information provided by the network layer to the data link entity on the neighboring node without flooding that node's storage capacity (it provides flow control). It also detects

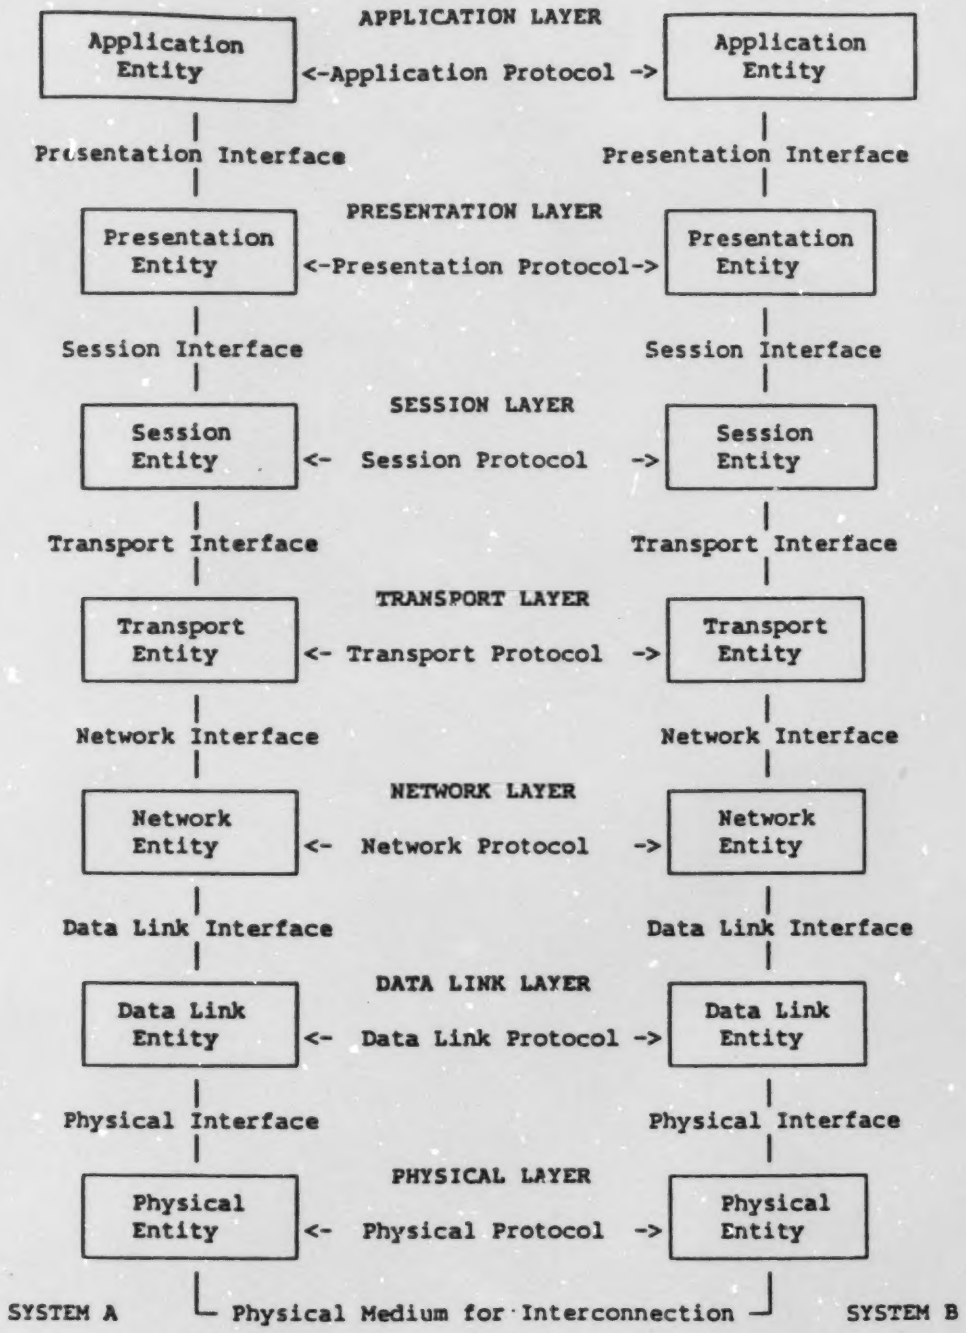

Figure 7. Open Systems Interconnection Reference Model.

transmission errors and corrects them or notifies the network layer that the connection is broken.

The physical layer provides the data-link layer with a transmission medium to the adjacent node; it defines the cables, connectors and signalling technique employed. Elements of the physical layer do not interpret the content of the data at all, including the boundaries between units of information.

This description ignores complexities introduced by the need for different classes of service and in the use of protocols designed before the present layered model was complete. These realities lead to some apparent redundancies between layers of the model. For example, end-to-end error checking may be offered by the transport layer, although error correction is performed between adjacent nodes by the data-link layer. If there are intermediate nodes in the subnet, an information unit can be received by a node that fails before forwarding the unit (a lost packet). Alternatively, erroneous receipt of an acknowledgement can lead the data-link layer to send information twice (a duplicate packet). Or, due to transmission delays, packet #1 may arrive at the destination before packet #2. But in a network with few transmission errors and a simple topology (where senders are adjacent to receivers) one can choose a data-link protocol that omits error correction, relying on the transport layer for this service.

As another example, transport services might be used to deliver digitally encoded voice messages over the same network that carries data. For voice transmission, speed is more important than quality since language carries enough redundant information to allow a listener to understand speech sent over an imperfect connection. Thus, the international transport protocol standard defines a class of service that does not include end-to-end error checking. Negotiation of such a class of service to deliver voice packets is fruitless, however, if matching services are unavailable from the network layer.

These examples show that adoption of standard protocols is not a panacea. One must still consider the applications that use communication services and the classes of service such applications require.

6.2 Long-Distance Transmission Media

To obtain transmission media for long-distance communication most individuals and institutions have two choices: lease or rent circuits from the telephone company or install earth stations and rent a channel of a satellite transponder. Both these basic services are provided by a heavily regulated company known as a common carrier. A major change in common carrier regulatory structure is about to occur, as any reader of the morning newspaper can verify. The effect of this action on service costs and availability is still a matter for speculation.

6.2.1 Telephone Company Services

6.2.1.1 The telephone network.

Because the network was established to transmit voice, we first examine it from the perspective of an ordinary telephone call. The telephone network comprises a six-level hierarchy. At the lowest level, the subscriber rents a "local loop" to the central office. To it, he connects a telephone set. Signals (an electric current) generated by the telephone traverse the local loop (a copper wire). After dialing, the switching system at the central office discovers whether the dialed number is found within the central office or outside it. If the number is outside the central office, the call is switched to "tandem lines" or "toll connecting trunks" that connect to a toll office. The toll office may be able to connect the call to the destination central office if it is nearby. Otherwise, switching continues up the hierarchy, over intertoll trunks, to primary, sectional and finally regional offices (and back down the hierarchy to the destination). A leased telephone line is similar, except that the route is permanently established.

The signal may thus travel over many transmission facilities before reaching its destination: as a current in copper wire, as light carried in an optical fiber, as microwaves radiated between two stations within line-of-sight, or as VHF radio signals radiated from a satellite earth station to a satellite transponder and back. But only from the sender's and receiver's premises to the central office does the call have exclusive use of the facility on which it travels. Across the tandem lines or the intertoll trunks the call is combined with many others to be transmitted together at high speed. In addition, the call almost always travels much farther than the original signal generated by the telephone reaches. Distortion and attenuation soon reduce the signal to inaudibility. Thus, at intervals, the signal must be amplified.[71]

6.2.1.2 Digital vs. analog signalling.

As anyone who has listened to a scratched phonograph record knows, when you amplify a signal superimposed by noise, the noise also grows louder. Taking an analog signal from a telephone and simply amplifying it repeatedly along its way leads to poor signal quality at the recipient's end. Fortunately there is a solution to this problem. If the analog signal is converted to a digital signal, it can be regenerated to its original state, zero or one, before it degenerates so badly that the regenerator can no longer tell the difference. This is what is done in central offices for over 80% of local loops.

---

71. John McNamara, Technical Aspects of Data Communication, p.94-95.

The coding-decoding function performed in the central office samples the analog signal and converts it to one of 128 values representing amplitude, plus one bit designating off-hook. At the receiver's end this digital value is again converted to an amplitude and transmitted to the listener's telephone. From Nyquist's signalling theorem we learn that our telephone line, which for historical reasons has a bandwidth of 4 kiloHertz, must be sampled at 8 kiloHertz. Thus our voice is being transmitted at 64,000 bits/second. If a computer could directly transmit digital information onto the subscriber loop, high speed data transmission would be available through very cheap interfaces at every telephone jack. In fact, such a digital transmission is available today by special arrangement (Bell's Digital Dataphone Service.)

By the end of the decade, all communication within the telephone network should be in digital form (Table 7). One stumbling block remains: to convert voice signals from ordinary telephone instruments into digital signals economically before they are placed on the local loop. Digital telephone signalling is already used in special phone sets connected to digital PBX's. What remains is to reach agreement on standards for the digital subscriber loop interface and to convert the large quantity of already-installed equipment to use the new technology.

6.2.1.3 Modems.

Until the network is fully converted, digital computer signals must be transformed into a form that expresses digital information within the 300 to 3000 Hz frequency range encompassing human speech. This audible signal is transmitted over the local loop, only to be interpreted there as a purely analog signal to be digitized for transmission through the phone system, there to be converted back to analog, and then to digital once more at the receiver's end. Imposing information on a signal of some chosen frequency is called modulation; regenerating the original information is called demodulation. The device used to perform this conversion at each end is called a modem. Information in its original, unmodulated form is called a baseband signal.

Table 7: Growth of digital facilities from 1980-1990.[72]

NETWORK COMPONENTS	'80	'90
Local loops	80%	90%
Local switching	40%	90%
Metro. transmission facils.	40%	90%
Toll switching	25%	92%
Local signalling	0%c	80%
Toll signalling	85%('85)	95%
Long distance transmission	1%	25%

cPrivate microwave and satellite links increase this figure.

Because the electrical parameters of an analog signal are easily distorted by side effects of the circuits that process it, analog circuits are tricky to build, requiring the application of many empirical design rules. If solving the digital statement of analog problems becomes tractable, as has occurred with the advent of high performance signal processing microcomputers, devices previously requiring costly analog circuitry will decline in price as digital techniques are applied.

Among the first devices to benefit by the application of digital signal processing techniques are modems. As an example, Bell 212 compatible modems designed for use with personal computers are now available for under $500. Such modems cost $700 to $1000 only two or three years ago. As of May there were eight low-speed modem chips available, with more on the drawing board.[73] Chips to do higher-speed modem functions as part of proprietary modems have also been described recently.[74] At present, old-technology 300 bits/second modems sell for $100 to $300. It is likely that similarly priced 1200 bits/second modems using these IC's will be available within two years. Only the smaller market place prevents higher-speed modems from being similarly reduced in price, but at least some decline is likely.

6.2.2 Value-added Networks

Because the dialed telephone network was not designed to transmit digital information, and because the rate structure makes lengthy long-distance calls expensive, institutions needing long-distance communication channels achieve higher bandwidth and greater reliability by leasing dedicated telephone

---

72. Falconer, W.E. and Skrzypczak, C. The Bell System - On Its Way to a Digital Network, Bell Laboratory Record 59, 5 (1981) p. 259-262.

73. Stephen J. Durham, Digital Design, May 1983, p.104-109.

74. Ibid., p.90-92.

circuits. Leasing a circuit to a distant city is also expensive, thus this alternative is not cost-effective unless the circuit receives high use. Management of leased circuits introduces operational complexity, since someone must order the circuit, monitor its quality and negotiate with telephone company technicians when problems arise. Many applications for long-distance data communication do not entail lengthy connections. Instead, periodic short interchanges occur. In such applications, one prefers a dialed network, but one designed to transmit high-speed data.

Since the early seventies, two such networks have been available, Telenet and Tymnet. These networks offer network, data-link and physical level services obeying an internationally recognized set of protocols designated as CCITT Recommendation X.25. CCITT is the international standards organization to which the telephone authorities of all countries may belong. X.25 provides a service interface to "virtual circuits", both leased and dialed. These are analogous to the physical circuits provided directly by the telephone company, complete with collect calls, procedures for placing calls, and disconnection.[75]

Data is placed on the network in units called "packets", leading to the term "packet switched data network" or PSDN. Because networks offering an X.25 service provide more than a simple transmission channel, they are also described as "value added networks". The provision of transmission services that add value will be deregulated under the new regulatory ground rules. American Information Services, "baby Bell", will offer a VAN service in competition with the established networks.

X.25 is not without defects. It offers a plethora of complex options, some implementing the same feature different ways. It redundantly supplies flow control at two levels. It does not fully specify the lowest three layers of the seven-level architecture, since only the protocols between the host (the "data terminal equipment") and the carrier (the "data communication equipment") are defined. What the carrier does with data between the time he receives it from its sender and delivers it to its destination is left unspecified. Despite these flaws, X.25 offers the only networking option implemented by a wide group of computer systems. It will continue to be the interface offered by carriers to long distance communication networks. Its careful use permits data interchange among many diverse systems.

VAN service charges are based on connection time and number of packets sent, contrasted with the distance-dependent leased line charge and the time and distance based direct dial charge. As a general comparison, VANs occupy a window between very short, infrequent connections (best served by direct dial) and lengthy, frequent connection (best served by leased lines).

---

75. A "datagram" service interface, analogous to sending telegrams, is optional under X.25, and U.S. networks don't offer it.

6.2.3 Satellite Transmission

Participants in satellite communication each install earth stations at their site, and together rent a portion of the capacity of a transponder (sender/receiver) installed on a satellite in geosynchronous (earth-stationary) orbit. Information sent to the satellite from any station (as high-frequency radio transmissions) are received by all the other stations roughly simultaneously about half a second later. As noted above, satellite transmission may even be used at intermediate stages in point-to-point transmission, since the public long-distance telephone network uses satellite links extensively.

6.2.3.1 Bandwidth.

Satellite transmission offers the potential of very high bandwidth communication. With a dedicated earth station, rates of 56 thousand up to 1.54 million bits/second are available. Lower bandwidth channels can be rented from satellite carriers without the need of a private ground station. Instead, a central ground station is shared among many users, and individual signals are carried to their destination over leased telephone circuits.

6.2.3.2 Reliability.

In telephone transmission, impulse noise is the most common source of transmission error. Impulse noise arises from many sources. Some examples include mechanical vibrations converted to electrical disturbances (noisy relays in a switching office); radiation from nearby electrical sources; electromagnetic pickup of a nearby signal. In space communication, however, the main sources of error are thermal and shot noise (from resistances through which the electric current passes in the antenna); fading, due to scattering of the signal off of obstructions in the atmosphere, principally rain; and interference from other signals, natural and man-made.[76]

One important man-made source of interference is from the signals arriving from other satellites. At present the FCC allows satellites in geosynchronous orbit to be spaced at four degree intervals as seen from the earth's surface. Under a pending FCC proposal, that spacing will be reduced to two degrees. Another man-made source of interference is terrestrial microwave transmission.

Ground stations overcome these sources of error in several ways. First, by making the "dish" size larger, the antenna sees a smaller angle of the sky. This minimizes the interference from other signals and increases the strength of the signal received. However, it makes the antenna more expensive; it makes the antenna heavier (thus making the placement of the antenna more

---

76. Mischa Schwartz, Information Transmission, Modulation and Noise.

difficult); and as the size of the antenna increases, its field of view decreases, ultimately requiring complex aiming mechanisms.[77]

Fading is overcome by engineering the satellite and ground station signal strengths to higher levels than are ordinarily needed. This is not difficult at the ground station, which is not severely limited in size or power consumption. The satellite transponder, however, is power limited, since it depends on photo-voltaic cells for energy. Fading problems are greater in the 12 and 14 GHz ("Ku") bands than in the 4 and 6 GHz bands used for most early satellite communication. Nonetheless, the Ku bands are attractive because they are less crowded.[78]

An interesting recent entry in the satellite communication arena, Equatorial Communications Company, sells small (.6 meter) low bandwidth (9600 bits/second) receive-only earth stations. This offering makes use of "diversity" techniques, long used to improve the error rate of military signals subject to jamming or weak telemetry signals from space probes. Essentially, the method used by Equatorial transmits a signal over many frequencies in the large bandwidth available to it. At the earth station, a microprocessor recombines the signal in such a way as to minimize the effects of expected types of noise and interference.

6.2.3.3 Operational complexity.

Because satellite communication employs radio transmission, the FCC requires that transmitting systems be monitored by personnel holding a commercial radio license. Such skills are not readily available; thus, most satellite services sell monitoring services along with their earth stations and transponder space.

6.2.3.4 Cost.

Satellite Business Systems, a consortium of IBM, Aetna Life and Casualty and Communications Satellite Corp., offers very high performance, very expensive satellite communication. 1.54 million bits/second channels are available through earth stations costing over a million dollars. The SBS earth stations provide combined voice, data and video transmission for large corporations.

6.2.4 Other Providers of Subscriber Connections

The only other communication medium that has geographical availability even approaching that of the telephone system and satellites is cable television (CATV). This facility is provided through a shared, high-bandwidth cable, tapped at every subscribing home in a neighborhood. At first shared use of

---

77. Roger L. Freeman, Telecommunication System Engineering.

78. Ibid.

this medium for transmission might seem to present a problem, but as we shall see in the next section there are convenient means allowing many users to share a single medium. Few cable operators, so far, have introduced such systems, despite much publicity.

This opportunity for wide access to data communication will probably not be generally exploited as soon as is technically feasible because so few services are available that might motivate a consumer to buy the termination equipment required. Introduction of such service also faces a regulatory morass. Perhaps by the end of the decade such applications will appear in residential communities. Before then, however, institutions that have private cable systems should certainly exploit them for data transmission.

6.3 Local Communication

In choosing local communication mechanisms, a planner is faced with a large number of alternatives. Just as in other communication technologies the choices boil down to ones of protocol. Commercial network offerings provide solutions for physical, data-link and network layers, but differ most significantly in their approach to physical layer problems. Topology, access protocols, signaling technique, transmission medium and connector design, form the physical layer in a local area network.

6.3.1 Network Topologies

When one computer system wishes to communicate with several others, one approach installs several wires, one pair between each system that wished to communicate. This configuration has the advantage that each pair of communicators transfers information at will, with no interference from any of the others. Its disadvantage is that for three partners there must be three pairs of wire; for four hosts, six; for ten, forty-five pairs; and so on, proportional to the square of the number of hosts. This proliferation of physical connections is both costly and an operational nightmare. Instead, real networks fall into one of four simpler classes. One important thing to note: the topologies discussed here are electrical. A transmission medium might be strung from one station to another in a variety of physical configurations, and still exhibit the same electrical "shape".

6.3.1.1 Star networks.

This topology designates one of the points in the network as a hub or central node. One connection runs from each of the other nodes to this hub forming a star. If two peripheral nodes wish to communicate, information must be routed through the hub. Star networks frequently connect terminals to central host computers. Their disadvantage when used to connect independent computer systems is the dependence placed on one point of the network. Should the central node fail, no one can communicate with anyone else. The RLIN network

is a form of star network. Because communication lines from the RLIN central computer go to more than one terminal, it is actually a tree; a polling protocol permits terminals on the same line to be connected to the host in turn. Since RLIN terminals currently have no independent processing capabilities it is not a disadvantage that communication must pass through the central computer. In a distributed system architecture, reliance on a single node reduces reliability.

6.3.1.2 Bus networks.

A bus network (conceptually identical to the I/O buses described in the chapter on processors) provide a single, high bandwidth roadway for data to travel. Each node of the network taps into the bus, and signals sent on the bus are heard by all stations. Transmission on bus networks requires a protocol for sharing the single data path. Receivers on bus networks play a passive role, watching for messages having their address. A passive receiver can easily be constructed not to disrupt the network if it fails or is disconnected, but detection of passive receiver outages is more difficult.

6.3.1.3 Ring networks.

As the name implies, data in a ring network passes from node to node. Receivers in a ring network are active. Should one fail, it disrupts transmission across the failed node. To prevent single point failures such networks often consist of two rings, operating in opposite directions. In single point failures, adjacent nodes notice the blockage and join the alternate rings, bypassing the node that failed.

```
...o-->--o-->--o-->--o-->--o... Original clockwise ring.
...o--<--o--<--o--<--o--<--o... Original counter-clockwise ring.

                A

...o-->--o->|--o--|>-o-->--o... Single ring after node A
...o--<--o-<|--o--|<-o--<--o... fails and is bypassed.
```

This permits quick identification of the bad node. As in the case of bus network topology, ring networks can be used in a contention mode and a contention free mode.

6.3.1.4 More complex topologies.

One can imagine building an arbitrarily complex network, occupying a position somewhere between a star and the n-squared connections of the initial example. By considering the seven layered distributed computer architecture it is easy to recognize that this complex connectivity introduces the need for a more complex routing protocol in the network layer. The bus, ring and star topologies are special cases that simplify the routing problem. Thus, these topologies are favored in current local area networks. Long distance packet

switched data networks exemplify an application for the more complex topologies.

6.3.2 Transmission Media

Having chosen a topology, some medium must be installed in the chosen configuration among all the stations. Of the host of possible media, two are common today and one will be increasingly important in future.

6.3.2.1 Twisted wire pairs.

This medium consists of two insulated copper wires, twisted together, perhaps shielded by an exterior foil wrapping, and surrounded by an insulating sleeve. Conservative estimates of a design using widely available RS-422 balanced interface drivers and receivers with 24 gauge twisted pair telephone wire predict acceptable operation at 1 million bits/second at distances up to about 400 feet.[79] Grounded shielding and thicker wire permit operation over longer distances or at greater speed, as does its use with a network design requiring active station taps. This medium is cheap and connector technology is readily available. Without shielding, twisted pair transmission at high data rates creates electro-magnetic interference that may affect radio/TV reception or operation of devices dependent on RF signalling, and may violate FCC regulations.

6.3.2.2 Coaxial cable.

Coaxial cable comprises an inner conductive core separated by an insulator from a grounded outer concentric cylindrical conductor. This medium, installed for community cable television systems, provides a total bandwidth of 300 MHz (50 6 MHz channels) with repeaters every ... feet. CCITT practice for large diameter (.375") coaxial cable operating at 12 MHz requires repeaters every three miles.[80] Many local area networks limit cable length to 1.5 miles at 10 Mbps so that cheaper connectors and less expensive transmitters and receivers can be used. Coaxial cable is an expensive medium, with a cost of $X/foot.

6.3.2.3 Optical fiber.

This medium offers bandwidth equalling coaxial cable with no electro-magnetic interference problems and a low cost per foot ($.xx). It has begun to be widely used in terrestrial long-distance communication. Its disadvantage for local communication is that connector technology has not equalled medium

---

79. Electronic Industries Association, "Electrical Characteristics of Balanced Voltage Digital Interface Circuits", EIA RS-422.

80. Freeman, op. cit., p.254.

production and transmission technology. Tapping an optical fiber cable leads to high signal attenuation, and splicing is time-consuming. Advances are being made, especially by researchers working for Bell Labs, and optical fiber should emerge as the communication medium of choice in two to three years.

6.3.3 Baseband vs. Broadband Local Area Networks

As noted in the discussion of modems above, baseband signalling merely refers to transmission of a signal in its original form, without transposing it to another frequency range (modulating it). Signal modulation serves two purposes: it permits frequency multiplexing of more than one signal on a high bandwidth (broadband) medium, and it matches the frequency spectrum of a single signal to the characteristics of the transmission medium. Thus, baseband and broadband are not opposite terms. In fact, if the frequency range of a baseband local area network is large, it is properly termed broadband baseband! As a descriptive term applied to local area networks, broadband implies modulation, usually for the purpose of frequency multiplexing on a medium capable of much greater bandwidth than that used by the LAN.

In the context of local area networks, baseband signalling implies lower cost per station juxtaposed against potentially shorter maximum distances and less effective use of the transmission medium. For short distance use of twisted pair, baseband signalling is an appropriate choice -- the medium's bandwidth is being fully exploited. In the case of long runs of coaxial cable, an institution might conclude that more effective use the medium was required to justify its cost or, indeed, to make its use feasible at all over cross-campus distances. A common application for entails transmission of a local area network signal within one or two standard channels on CATV coax.

Baseband and so-called broadband techniques do not conflict. An increasingly common design approach employs baseband signalling within geographically close subdivisions of a local network (for example, one floor within a building). Then, using devices called bridges, these subnetworks communicate over a very high bandwidth medium, using modulation to achieve longer distances and to allow sharing of the medium with other applications (for example, video). The bridges listen to traffic on the subnetwork and forward messages destined for other subnetworks. Thus, only the bridges require higher cost transmission hardware, and since the most complex cabling occurs between stations within buildings, a low cost per foot and cost per station results.

6.3.4 Access Protocols

Access protocols fall into three categories: centrally controlled protocols, distributed collision-free protocols and contention protocols. Because centrally controlled protocols rely on the robustness of a favored "master" node, local area network designers prefer the other protocols. These can be contrasted by their performance in initiating individual transactions (contention protocols excel here) and their behavior under load (collision-free protocols achieve better peak-load performance).

6.3.4.1 Contention protocols.

Contention protocols take advantage of the fact that independent host systems connected to a high speed medium do not transmit data very often or very long. If this is true, then on the average, when a system wishes to send a block of data the bus is likely to be unused. Under contention protocols, a system checks to see if a message is on the bus by sensing for the presence of a carrier signal. If none is seen, the system transmits its data block. This protocol is called "carrier sense multiple access - collision avoidance" or CSMA-CA.

One problem arises with this technique. Transmissions take a small amount of time to traverse the bus. If two systems begin to transmit messages at approximately the same time, neither will see a carrier present on the bus, yet a collision will occur. Data-link protocols require that messages have check sums enabling the receiver to verify correct transmission and establish acknowledgement procedures and timeouts notifying the sender that a message was or wasn't received. CSMA-CA systems rely on this error-detection and acknowledgement procedure to recover from the "simultaneous transmission" collision. OmniNet, a microcomputer local area network offered by Corvus Systems, uses this technique.[81]

In another contention protocol, "carrier sense multiple access - collision detection" or CSMA-CD, the device used to transmit data blocks simultaneously receives data from the bus and compares it with the data transmitted. Should a collision occur, it is detected almost immediately, and an abort sequence is transmitted. This technique, exemplified by Ethernet, is more complex, and thus more costly than CSMA-CA, but offers more effective use of the transmission medium.[82] Recently, VLSI implementations of the Ethernet physical protocol have become available. This will eventually reduce the cost from over $1000 per station to a more attractive $200 a station. Intel, Ungermann-Bass/Fujitsu, and Seeq offer Ethernet chip sets. Seeq's chip was developed in concert with 3Com Corporation, whose president, Robert Metcalfe, is one of the patent holders for the Ethernet technique. CSMA/CD (the Ethernet access protocol) has received final approval as part of an IEEE local networking standard (IEEE P802.3-83/.01).[83]

The two access protocols just described depend on all stations being able to "hear" if any one station transmits, that is, they operate on bus topologies. Ring networks also can use a contention protocol. A station on a ring can transmit whenever the ring is idle. If it receives someone else's packet

---

81. William G. Wong, "An Introduction to Local Area Networks: Part II", Microsystems, November 1983.

82. Andrew S. Tanenbaum, Computer Networks, p.292-295.

83. A. Lyman Chapin, "Computer Communication Standards", Computer Communication Review, July 1983, p.38.

while it is transmitting, it holds the packet in a register and retransmits this packet when it is through. This technique, called register insertion, is used in the IBM Local Communications Controller for the Series 1.[84]

Contention protocols face a common problem: when the transmission medium is lightly loaded they work well, but as activity rises collisions and retransmissions occur more frequently. Key to the effective use of contention protocols under load is the algorithm chosen for selecting the time of retransmission. Even when this time is selected according to algorithms that dynamically adapt to system load, performance of a collision network degrades when the cable is long and the packets are short.[85]

6.3.4.2 Distributed collision-free protocols.

Ring networks circulate a special bit pattern when the network is idle. A station wishing to transmit waits until this pattern (the "token") is in its receiver, sends a modified form of the token called a connector, then transmits its message. Bus networks can use a modified form of this protocol that is sometimes called a logical ring because it entails explicitly passing control of the bus to the "next" node.

Since no bandwidth is lost to collisions token passing protocols exhibit very good behavior under heavy load. With their active connections token passing access methods also have the advantage of reducing distance restrictions. This advantage, as usual, has a flip side: when many nodes are connected to a token passing network the stations are subject to a hardware problem in which they may "lose their place" in decoding the bit stream. Another problem arises if an error causes the token to be lost.

IBM has described a token-ring network developed at their Zurich lab. This network employs an active token monitor node to perform recovery of lost tokens and to log other system faults. This network is described as a "star ring", because it is intended to be installed using special cable running from stations on the network to nearby wiring concentrators. The cable preserves an electrical ring topology, but physically the network comprises many stars, interconnected by a higher bandwidth backbone cable. This installation design is intended to mimic telephone cabling in a building, making maintenance easier.[86] The Nestar Plan 4000 network has a similar cabling strategy but uses a bus electrical topology.

Token ring and token bus protocol standards are under consideration before the IEEE 802 committee, but are not yet final.

---

84. R.C. Dixon, et.al., "A Token-Ring for Local Data Communications", IBM Systems Journal 22 #1-2, 1983, p.47-62.

85. Tanenbaum, op. cit., p. 296.

86. Dixon, op. cit.

6.3.5 Commercial Local Systems

6.3.5.1 Baseband LAN's

Table 8 lists several commercially available baseband LANs, showing the medium, protocol and electrical topology each employs. Table 9 shows station and file server prices for a few of these systems. As personal computers and workstations proliferate, more LAN products becaome available. These represent the most cost-effective way to share expensive peripherals within an office or building, and make possible information sharing as well. However, it is important to note that while many systems offer a communications subnetwork, few offer complete support for shared files, print spooling, remote submission of background jobs to large mainframe computers, electronic mail, and other applications. In terms of the seven-layered architecture, only the first three layers are now widely available.

Table 8: A Few Baseband LAN Implementations[87]

NETWORK	Mbps	MEDIUM	PROTOCOL	TOPOLOGY	MAXIMUM NODES	MAXIMUM METERS
Nestar Systems Cluster/One	0.24	Wire	CSMA/CA	Bus	64	305
Corvus OmniNet	1.00	Wire	CSMA/CA	Bus	64	1200
North Star NorthNet	1.00	Wire	CSMA/CA	Bus	64	1200
Zilog Z-Net	0.80	Coax	CSMA/CD	Bus	255	2000
Intel and NCR Mirlan	1.00	Coax	CSMA/CD	Bus	?	1200
Apple Computer AppleNet	1.00	Coax	CSMA/CD	Bus	128	600
Orchid,AST,SantaClaraSys.PCnet	1.00	Coax	CSMA/CD	Bus	10	?
Novell Sharenet X	1.40	Coax	CSMA/CD	Bus	10	?
Xerox Ethernet	10.00	Coax	CSMA/CD	Bus	1024	2500
3COM Ethernet	10.00	Coax	CSMA/CD	Bus	1024	2500
DEC Ethernet	10.00	Coax	CSMA/CD	Bus	1024	2500
Intel Ethernet	10.00	Coax	CSMA/CD	Bus	1024	2500
Ungerman/Bass Ethernet	10.00	Coax	CSMA/CD	Bus	1024	2500
Network Systems HYPERbus	10.00	Coax	CSMA/CD	Bus	128	730
Lantech uNETix	10.00	Coax	CSMA/CD	Bus	128	600
Network Systems HYPERchannel	50.00	Coax	CSMA/CD	Bus	16	1500
Gould Modway	1.54	Coax	Token	Bus	250	4500
Datapoint ARC	2.50	Coax	Token	Bus	255	610
Datapoint ARC	2.50	Fiber	Token	Bus	255	1600
Standard Microsystems ARC-PC	2.50	Coax	Token	Bus	255	?
Nestar Systems PLAN 4000	2.50	Coax	Token	Bus	255	6706
Northern Telecom Omnilink	0.04	Coax	Token	Ring	20	1500
Vector Graphics LINC SABRE-Net	0.75	Wire	Token	Ring	16	3200
Prime Computers Primenet	10.00	Coax	Token	Ring	247	3048
Proteon Associates proNET	10.00	Coax	Token	Ring	255	2400
Logica, Inc. Polynet	10.00	Coax	Token	Ring	?	?
Ztel Inc. Pvt.Netwk.Exchg	10.00	Coax	Token	Ring	?	?
Apollo Computer Domain	12.00	Coax	Token	Ring	?	900

---

87. William S. Wong, op.cit.

Table 9: LAN Station and Server Prices

NETWORK	PRICE NODE	SERVER
Novell Sharenet S	$ 250	$ 8795
Nestar Systems Cluster/One	$ 400	
Corvus OmniNet	$ 495	$ 3985
Standard Microsystems ARC-PC	$ 495	
Ungerman/Bass Ethernet	$ 500	
Novell Sharenet X	$ 595	$ 4995
Nestar Systems PLAN 4000	$ 595	$19995
Destek Group Desnet	$ 600	
Orchid,AST,SantaClar PCnet	$ 695	
3COM Ethernet	$ 950	$12500

6.3.5.2 Broadband LAN's

Many fewer broadband systems are available, reflecting the fact that such systems have a more limited range of applications. Examples of applications where the extra cost and complexity of broadband signalling is justified follow:

- **Connecting buildings on campuses.** Distance limitations of baseband networks can be overcome through the use of modem technology.

- **Using existing cable television facilities to implement local area networks.** Broadband networks frequently multiplex a baseband signal on one or several TV channels in the standard channel allocation used by commercial CATV.

- **Sharing a high bandwidth medium with other applications.** If an institution needs to use single medium to transmit voice, fax, data and video, a frequency multiplexed broadband medium may be an effective choice.

Commercial broadband systems include Sytek, a CATV product; Ungermann-Bass, Ethernet linked with a broadband backbone; and Wangnet, from the maker of word processors and office automation, a broadband system. NEC offers a family of networks, including broadband coax and optical fiber systems that accomodate video, fax and voice along with data.

6.3.5.3 Digital PABX's

Unlike the local networks previously described, digital private automatic branch exchanges employ telephone wiring installed in a centrally controlled star topology as switched data circuits. Because PCM encoding provides a 64 Kbps bandwidth between the telephone and the exchange, applications seeking intermittant access to a few central resource are well served by such an

approach. One wiring installation may serve both voice and data needs. Often, a special telephone set provides an RS232-C interface and permits data transfer at up to 9600 bps simultaneously with voice traffic.

Using a digital PABX for data traffic may also present problems, however. Digital PABXs are designed internally to multiplex PCM encoded voice over an internal bus. Thus, the 64K bandwidth that is conceptually available to each individual connection may not be available to a large number of connections simultaneously. Call setup time (the time to switch to a resource connected to the PABX) is long compared with the time required to begin transmitting data across a LAN.

Perhaps the best use for digital PABX data communication is as a terminal gateway to a local network designed for data traffic or to a remote packet switched network. This permits relatively low speed terminal traffic to be multiplexed by the PABX control processor. Characters from several terminals can then be sent as a single packet over the data network to a central host. PABXs designed to work in concert with local area or long-haul data networks include the ROLM CBX, the Datapoint ISX, the Northern Telecom SL-100, and the InteCom Inc. IBX.

6.4 Cost Trends

Local communication presents more options to network planners than can reasonably be investigated. The availability of network options will increase, but some narrowing in their variety should occur as standards are adopted. Price per station connection is converging on $500 or less. Shielded twisted pair is a more cost effective choice for a transmission medium than coaxial cable, especially as optical fiber will become practical for higher bandwidth applications in the next two to three years.

Competitive uncertainty clouds the long-distance communication picture. If deregulation has the effect predicted by many analysts, long-distance rates will decline and competition in the value-added network arena will increase, lowering prices and adding function. The evolution of an integrated digital services network increases the likelihood of lowered long-distance data transmission cost, as does long-distance capacity increase due to installation of fiber-optic cables.

The likely decrease in terrestrial long-distance prices together with growing congestion of satellite frequencies makes satellite communication look less favorable as a long term choice for smaller communication users. Satellite communication may be an appropriate choice where an immediate application will quickly pay back the initial investment.

Chapter 7

Software Engineering

A trend noted in earlier chapters is that the cost of computer hardware has declined spectacularly over the history of the industry, and this is expected to continue for the forseeable future. As hardware costs decrease, software becomes a larger component of total system costs. Furthermore, the inherent reliability of solid-state electronics coupled with fault-tolerant hardware architectures leaves software reliability as the key factor in overall system reliability. For these two reasons, a great deal of attention has been focused in the last decade on "software engineering," the processes by which computer software is designed and implemented. This section will present an overview of the key concepts and innovations in this field, and attempt to project where these trends are likely to lead in the next decade.

The emphasis on engineering is of fundamental significance. Until recently, software development was considered to be an art, more akin to writing a novel than to designing a building (or computer hardware). In contrast, (hardware) engineering is considered to be a professional discipline, with well defined standards and methodologies that are codified in textbooks and taught in classrooms. Building such discipline in software development is a natural step in the maturation of the field of computing.

7.1 Software Quality

Perhaps the most important goal of software engineering is maximizing software quality. This is not a simple one-dimensional characteristic of a program; instead, quality has many interrelated aspects. Boehm, et al., have provided a very useful taxonomy of the various aspects of software quality, which is summarized in Figure 8.

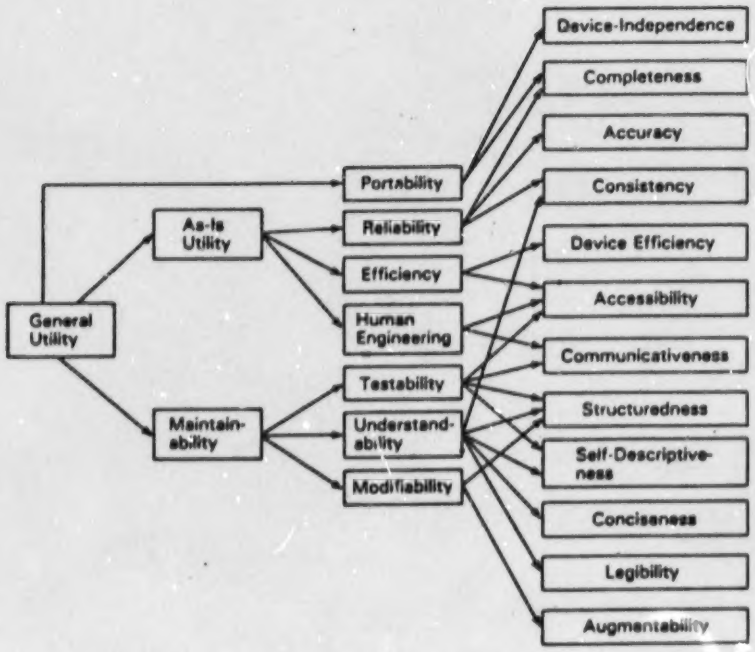

Figure 8. Software Quality Characteristics Tree [88]

The terms used in Figure 8 may be defined as follows: [89]

- Accessibility: Accessible code facilitates selective use of its parts.
- Accuracy: Accurate outputs are sufficiently precise to satisfy their intended use. ...
- Augmentability: Code is augmentable to the extent that it can easily accommodate expansion in component computational functions or data storage requirements.

---

88. Boehm, B.W., et al. Characteristics of Software Quality. Amsterdam and New York: North-Holland, 1978. Reprinted with permission.

89. Aron, J.D. The Program Development Process, Part II. Reading, MA: Addison-Wesley, 1983. p. 310-312. Reprinted with permission.

- Communicativeness: Communicative code facilitates the specification of inputs and provides outputs whose form and content are easy to assimilate and are useful.

- Completeness: Code is complete when all its parts are present and each part if fully developed. This implies that external references are available and required functions are coded and present as designed, etc. Completeness implies self-containability and robustness.

- Conciseness: Concise code contains no excessive information. This implies that programs are not excessively fragmented into modules, overlays, functions, and subroutines nor that the same sequence of code is repeated in numerous places, rather than defining a subroutine or macro, etc.

- Consistency: Internal consistency implies uniform notation, terminology, and symbology within the code; external consistency means that the content is traceable to the requirements.

- Device-Independence: Device independent code can be executed on computer hardware units other than the current one; e.g., any tape or disk that is available can be used to store a file.

- Efficiency: Efficient code fulfills its purpose without waste of resources. This implies that choices of source code constructions are made in order to produce the minimum number of words of object code; or that where alternate algorithms are available, those taking the least time are chosen; or that information-packing density in core is high. Device-efficiency means that each device or physical resource is used as intended; i.e., its advantageous features are optimized consistent with its role in the system.

- Human-Engineering: Code is human-engineered to the extent that it fulfills its purpose without wasting the users' time and energy, or degrading their morale, or misleading them.

- Legibility: The function of legible code is easily discerned by reading the code. ...

- Maintainability: Maintainable code facilitates updating to satisfy new requirements or the correct deficiencies. This implies that comments are used to locate subroutine calls and entry points, visual search for locations of branching statements are their targets is facilitated by special formats, or the program is designed to fit into available resources with plenty of margins to avoid major redesign, etc.

- Modifiability: Code is modifiable to the extent that it facilitates the incorporation of changes, once the nature of the desired change has been determined. Note the higher level of abstractness of this characteristic as compared with augmentability.

- Portability: Portable code can be operated easily and well on computer configurations other than its current one. This implies that special

language features, not easily available at other facilities, are not used; or that standard library functions and subroutines are selected for universal applicability, etc. ...

- Reliability: Reliable code can be expected to perform its intended functions satisfactorily. This implies that the program will compile, load, and execute, and that the program will continue to operate correctly, except for a tolerably small number of instances, while in operational use.

- Self-Descriptiveness: Code is self-descriptive to the extent that it contains enough information for a reader to determine or verify its objectives, assumptions, constraints, inputs, outputs, components, and revision status.

- Structuredness: Structured code has a definite pattern of organization. This implies that evolution of the program design has proceeded in an orderly and systematic manner and that standard control structures have been followed in coding the program, etc.

- Testability: Testable code facilitates the establishment of verification criteria and supports evaluation of its performance. This implies that requirements are matched to specific modules or diagnostic capabilities are provided, etc.

- Understandability: Understandable code is clear to the inspector. This implies that variable names or symbols are used consistently, modules of code are self-descriptive, and the control structure is simple or in accordance with a prescribed standard, etc.

- Useability (As-Is-Utility): Code is usable to the extent that it is reliable, efficient, and properly human-engineered. This implies that the function performed by the program is useful elsewhere, is robust against human errors ... or does not require excessive core memory, etc.

The arrows in Figure 8 denote dependency relationships between characteristics. For example, the general (long-term) utility of a program depends on its being immediately useful (As-Is-Utility), maintainable, and portable to possible future hardware environments. Likewise, As-Is-Utility requires the program to be reliable, efficient, and human-engineered.

7.2 Software Life-Cycle

A very useful framework for a discussion of software engineering principles is the concept of the software life-cycle. This term refers to the entire history of a piece of computer software: requirements definition; external and internal design; implementation and testing; and ongoing support, maintenance, and enhancement once the software is in use. The software life-cycle is shown schematically in Figure 9.

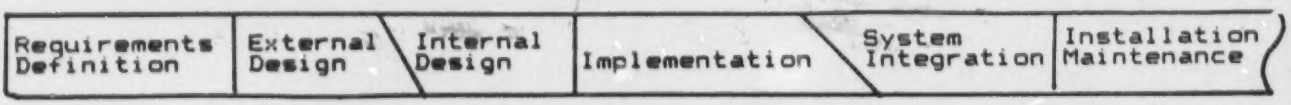

Figure 9. Software Life-Cycle

The tangible product of each stage of design is a document that records all decisions made (and the reasons for them) and serves as input to the next stage of the process. Of course, the real design process is never quite this clear-cut. The boundaries between the various design stages are not absolute; indeed, there is not always even a sharp dividing line between design and implementation. Instead, there is usually some iteration between successive stages of the development process. For example, a difficulty or new idea encountered in the internal design process may suggest or require a change in the external design or even in the requirements. Likewise, changes in the internal design may occur during implementation. Because of this, it is critical to institute change control procedures at the outset. These specify the review procedures and sign-off required to make retroactive design changes and the procedures for keeping the relevant design documents up to date.

One final comment on software development methodology should be made. While in most cases the controlled, orderly approach to development characterized by a formal software life-cycle is the best way to proceed, there are cases in which informal prototyping is more appropriate. This occurs when requirements are not well understood and some experimentation with one or more prototype systems is required in order to refine them. For example, alternative user interfaces might be mocked up in order to evaluate their relative effectiveness. Very high level declarative languages or application generators (see below) can be very effective tools for building prototype systems, especially as efficiency is seldom a concern at this stage.

7.2.1 Requirements Definition

Before a software system can be designed, the problem that it is to solve must be understood and stated clearly, along with any constraints and requirements on performance, development schedule, and cost. It is at this stage that the appropriateness of a computer-based solution to the problem should be considered explicitly. This requirements definition step seems obvious, but is often not given the attention it deserves. It should be kept as separate as possible from the subsequent step of external design, to maximize the chances that all requirements that can be identified at this stage are made visible and incorporated into the design.

7.2.2 External Design

External design (also called the functional specification phase) is the process of specifying how the system will appear to its users: the inputs and outputs, the transformations applied to the inputs to yield the outputs (but not the details of how the transformations are performed), error conditions and their handling, etc. If the system being designed is database-oriented, the database must be designed at this point if it does not already exist.

A particularly important aspect of the external design of online systems is the choice of user interface style. In times past, the initial choice was between a line-oriented and screen-oriented interface, but given modern display terminal technology, especially when it is coupled with a highly intelligent workstation, it is hard to imagine new systems being designed with line-oriented user interfaces. Having chosen a screen-oriented approach, however, there are still several alternative kinds of interfaces possible. The major distinction is between command-driven and menu-driven systems. In the former, the user must remember and type commands to instruct the system as to the next operation desired; in some systems, the syntax of such commands can be very complex. In the latter, the user selects the operation desired by picking an item from a list of possibilities that is displayed by the system. These menus may be hierarchical, i.e., a choice may lead to the display of another lower level menu from which varients of the operation may be chosen, etc. There are two methods in use for choosing menu items: The item may be selected by typing a number or letter, or by pressing a specified function key. Alternatively, the item may be selected by means of a pointing device, e.g., a light pen, a touch screen, or a "mouse." Such pointing interfaces are rapidly gaining popularity, especially in conjunction with the use of high-resolution raster graphics displays.

An area of software technology that will have a profound effect on user interface design when it reaches maturity is natural language understanding. To date, demonstration systems can understand statements only within a limited domain of discourse represented in their "knowledge bases," and this is likely to remain the case for the forseeable future. However, the techniques are well enough developed to be considered as a user interface for certain database systems where ad hoc queries from untrained users are to be supported, and the semantics of meaningful queries are well enough

circumscribed.

Voice input has recently moved out of the laboratory and is beginning to find its way into some products. Current products can recognize only a limited repertoire (on the order of 100 words) and must be taught to understand a particular speaker by repetition and correction, but the combination of high speed signal processing chips and more sophisticated natural language recognition software will certainly expand the range of applicability of this technique. It is too early to make any definitive statement on its ultimate utility. Voice output, on the other hand, has been developed to the point that good speech synthesis hardware and software is readily available for use in designing systems.

7.2.3 Internal Design

Once the external design of a new system has been specified, its internal design must be determined. Devoting sufficient time and resources to this phase of development is critical to the success of the project. Various studies have shown that around 70% of all errors in a large and complex system [90]

The basic problem of internal design is defining an optimal module structure for the software. A number of formal strategies have been developed to aid the designer in developing a system design that minimizes both development and maintenance costs. [91] These strategies revolve around the empirical observation that the best systems are those characterized by low coupling and high cohesion. Coupling refers to the interaction between modules of a system; the best system designs are those in which the modules have very narrow interfaces, i.e., minimal and controlled interactions. Cohesion refers to the rationale by which functions are grouped into a single module; the best designs are those in which modules perform only single well defined functions. Furthermore, good designs are usually well factored; that is, decision-making is concentrated high in the module hierarchy, while processing tends to be carried out lower in the hierarchy. Finally, good designs usually are characterized by high fan-out in the high level modules (but not too high, i.e., <= 7 other modules called) and high fan-in at the lowest levels.

It is impossible to overemphasize the importance of developing a coherent overall architectural framework for a complex system before detailed internal design is undertaken. One approach that has proven to be significant value is layering. A design is said to be layered if the modules it comprises may be grouped (into layers) in such a way that modules in any layer use only the functions and services provided by modules in their layer or by those of the

---

90. Zelkowitz, M.V. Perspectives on Software Engineering. Computing Surveys 10, 2 (1978), p. 200.

91. Yourdon, E., and Constantine, L.L. Structured Design. Englewood Cliffs, NJ: Prentice-Hall, 1979.

next lower level. Thus, a layer implements an abstraction that hides all lower layers. For example, a database management system that presents a very high level abstract view of data (e.g., relations) to its users might be built on the services provided by a record manager, which in turn is built on the services provided by a primitive file system. Another well known example is the ISO Open Systems Interconnection Reference Model, which defines a layered architecture for computer networking.

Operating systems are nearly always designed with a layered architecture. [92] A particularly interesting example of layering in an operating system is the virtual machine concept as embodied, for example, in the IBM VM/SP product. In such a system, a virtual machine monitor (CP in VM/SP) constitutes a well defined layer that presents a virtual machine abstraction to the layers above it. It multiplexes the hardware in such a way that multiple concurrently executing virtual machines may be supported by a single physical machine; these virtual machines appear identical to the physical machine except for some of their performance characteristics. Thus, it is possible to run multiple operating systems on the same hardware, or to provide an interactive timesharing environment by running multiple copies of a single user operating system (CMS in VM/SP). In contrast, most multiprogramming operating systems multiplex the hardware at a higher level, providing abstractions like "processes" or "tasks" that may execute concurrently. [93]

An additional method of structuring large systems is possible in an operating system environment that makes suitable provisions for it -- structuring the system as a collection of requestors and servers. A server is a process that offers a service (e.g., database access) to any other process (a requestor) that is authorized to use the service. This can be a very effective means of modularizing a large system because it enforces a high degree of isolation between functions; processes can communicate only via interprocess messages. To make this approach possible, the operating system must provide an appropriate interprocess communication mechanism that includes suitable access control facilities.

7.2.4 Implementation

Just as the design of software can be approached in a structured way, so can its implementation. Structured programming is the name given to a collection of programming techniques whose objectives are greater program readability and clarity, and a reduction in the time spent testing and debugging new code, resulting in greater programmer productivity. Experience with these

---

92. Peterson, J., and Silberschatz, A. Operating System Concepts. Reading, MA: Addison-Wesley, 1983.

93. It is amusing that such a system may be run on a virtual machine, further dividing the resources available to the virtual machine over multiple processes or tasks. However, unless explicit provision is made for cooperation between the virtual machine monitor and the "guest" operating system, efficiency can be substantially reduced.

techniques has demonstrated their efficacy; 5-fold increases in programmer productivity are not uncommon. [94]

One major theme of structured programming is the elimination of the GO-TO statement, because its use leads all too easily to unstructured [95] code that is error-prone, difficult to test, and all but impossible to maintain and extend. Languages that support structured programming provide well-structured looping (DO-WHILE) and branching (IF-THEN-ELSE) constructs in place of (or more usually, in addition to) the GO-TO statement. These constructs have only one entry and one exit, and hence any program (which itself may be considered to have one entry and one exit) can be composed by concatenating a series of them. Such a program is much easier to read and understand than its unstructured equivalent.

When a program is built using the structured programming constructs, a top-down approach to implementation is natural. This process is also known as stepwise refinement: First the overall control and input/output structure of the program is defined, freely using placeholder comments (called stubs) to represent functions whose details need not be considered at that level of the program. Then the stubs are expanded into code, possibly introducing lower level stubs in the process. This process is carried out iteratively until code has been written to perform all necessary functions. A program that has been written in this top-down fashion is generally easy to read and understand. More important, it is much less likely that its developer(s) will have failed to account for error conditions, end cases, and the like. The neglect of such special cases is a major cause of latent bugs in a program, i.e., those that show up (long) after the program has been placed in production.

One of the original motivations for the development of the techniques of structured programming was to make it easier to construct formal proofs of correctness for computer programs. Even with the resulting simplifications, however, this is a very difficult problem, and such proofs (even those a mathematician would characterize as informal) have been carried out only for relatively simple algorithms. It is not unusual that a proof is longer than the program to which it applies! [96] Nevertheless, work in this field continues, and it is to be hoped that eventually techniques will be developed that can be applied to the validation of real-life system.

---

94. Yourdon, E. Techniques of Program Structure and Design. Englewood Cliffs, NJ: Prentice-Hall, 1975. p. 143.

95. The terms "spaghetti" and "rats nest" are often used to describe this type of program.

96. Yourdon, E. ibid. p. 142.

7.2.5 System Integration and Testing

One question that remains, even if the top-down programming technique is used to implement individual modules, is whether the design as a whole should be implemented and tested top-down, bottom-up, or by some intermediate strategy. The traditional approach to building a system has been to implement the lowest level modules first and test them extensively using so called "test drivers" written especially for the purpose. As higher level modules of the system are implemented, they are integrated with the already existing lower level modules and tested in turn. Ultimately, the entire system has been built and can be tested as a whole. In a large project, with many programmers and much opportunity for misunderstanding and miscommunication, this system integration testing can become a disaster, in which all the previously tested modules or subsystems that were thought to work are discovered not to work together as a result of interface mismatches.

Top-down programming suggests another approach to system integration, namely, implementing the modules of the system starting with those at the highest level and proceeding downward. Rather than writing test drivers, dummy routines are written that simulate as-yet-unwritten routines at the next lower level. (One assumes that these dummies are substantially easier to write than the routines that they simulate.) Thus, system integration is carried out over the entire course of the implementation effort, and there is no chance of an eleventh hour disaster. An additional advantage of the top-down approach is that a working system (or the appearance of one) exists early, bolstering both the morale of the development team and the confidence of the project's management/sponsors/customers.

Very large projects, especially those that can be divided into layers as discussed above, may be carried out best by parallel development of the several layers. Integration of the layers may proceed in a top-down or bottom-up direction, depending on circumstances. However, this approach requires very careful specification of the interfaces between layers, and diligent project management to ensure that those interfaces are understood and adhered to.

As a large system undergoes integration testing (and as maintenance is performed after it is put into production), it is important to define procedures to ensure that the software continues to pass all previously successful tests. The approach usually taken is to organize a comprehensive set of tests (expanded as necessary) into a regression test that is run on each new version of the system. It is usually desirable to automate the comparison of the output of a regression test to its output when run on the preceding version of the system.

7.3 Development Tools

A great deal of attention has been focused in recent years on the program development environment: that set of tools used by a programmer in implementing, testing, and debugging software. Such tools fall into two major classes: 1) those that are used only as adjuncts to the process of development, e.g., program editors, debuggers, source control systems; and 2) those that are incorporated into the final product, e.g., parsers, macro processors, subroutine libraries. Both can increase programmer productivity significantly, but the following discussion will address only the first class.

7.3.1 Programming Languages

The programming language is clearly a programmer's most important tool. By the same token, the choice of programming language is one of the most important decisions made in the course of a development project. Languages may be grouped into three classes: low, intermediate, and high level.

Assembly language and most system implementation languages (SILs) are low level languages, the objects manipulated in the language are not very far removed from the underlying hardware. These languages give the programmer maximum flexibility and control of the hardware, but it is very easy to make programming errors that are difficult to locate and correct in testing. Debugging and maintaining a low level language program tends to be very expensive. Because these languages are hardware-specific, programs written in them are for the most part not portable.

High level languages (e.g., Pascal and Ada) allow the programmer to define and manipulate abstract data types (integers, reals, arrays, records, etc.) and strongly enforce the typing by ensuring that the operations carried out on a data object of a given type are meaningful for that type. In addition, these languages have been designed to support structured programming and the clean modularization of programs. Strong typing and the extensive compile time checking that it allows eliminate many sources of error and make debugging and maintaining a high level language program much more efficient than would be the case with a low level langugage implementation of the same program. Because high level languages carefully isolate the programmer from the underlying hardware and operating system environment, programs written in them are highly portable. For almost all new software development, a high level langugage is by far the best choice today.

In our classification scheme, intermediate level languages are those that were considered to be high level 10-20 years ago. Examples are FORTRAN and COBOL; PL/I may be classified as either an intermediate or a high level language, depending on how it is used by the programmer. Some SILs may be considered to be intermediate level languages, also. These languages manipulate typed data,

but the typing is not strongly enforced. They generally lack one or more of the constructs and facilities that make structured programming possible or practical. Because these languages have existed for some time and both are standardized and widely implemented, programs written (carefully) in them can be portable.

7.3.2 Program Editors

A program editor is used to prepare program source text interactively. While any editor will serve, it is possible to build some additional intelligence into the editor that makes it easier to use for program development. An example is a mode in which the cursor is automatically indented on a new line so that it is aligned with the first non-blank character of the preceding line. This makes it easy to enter the statements of a language like Pascal or Ada in which a nested block structure is typically accentuated by indentation.

A further advance is the development of syntax-directed editors. These have sufficient knowledge of the programming language to prevent the entry of syntactically incorrect programs. This is generally done by making the objects on which the editor acts those defined by the language syntax rather than characters, words, sentences, etc., as would be appropriate for a text editor. Without such an editor, at least one compiler run (and often more) is required just to eliminate syntax errors from a program before testing can even begin. Thus, a syntax-directed editor can increase a programmer's efficiency. More important, with a properly designed user interface, such an editor actually can aid the programmer in developing the program in a top-down structured manner.

7.3.3 Symbolic Debuggers

Significant productivity gains are also possible in the debugging process through use of a symbolic debugger. Early interactive debuggers operated at a machine-language level. To use them to debug a high level language program, one has to understand in detail how the compiler maps program statements into machine code; moving back and forth between the source language and machine code is generally difficult and time consuming. This was unquestionably an advance over dump reading, but a symbolic debugger that interacts with the programmer in terms of the (high-level) constructs of the language being used is a much better tool. With such a debugger, it is possible to read or modify the value of a variable that is identified only by name, or to set a breakpoint at a statement label or procedure entry point, again specifying only its symbolic name.

An even more powerful tool can result from the fusion of a program editor, and incremental compiler or interpreter (so that a separate compilation step is not needed), and a symbolic debugger. Then it is very easy to move freely between program editing and debugging. For example, a runtime error results in a return to the editor, with the cursor pointing to the statement where the error occurred. The error then can be corrected and the program retried, all without leaving the integrated program development environment.

7.3.4 Source Control Utilities

When large systems are implemented by programming teams, keeping track of source code and controlling its integration can become overwhelming if carried out by manual procedures. Fortunately, computers are very good at such clerical tasks, and hence source control utilities have been developed. Such a utility typically stores all versions of a module along with descriptions of the changes made from one version to the next; this makes it easy to return to an earlier version when necessary and provides an audit trail of changes. Another function of a source control utility is to automate the system integration process (e.g., linkage editing) to ensure that each test version of a developing system is properly constituted. It might also be used to keep track of test cases, test data, etc., and to automate the testing process. Once the system has been placed into production, a source control utility also can be used to support maintenance activities, keeping track of both test and production versions of the system and controlling the testing and introduction of new versions.

7.3.5 Development Support Hardware

Mainframe-based development often suffers from a combination of non-optimal tools and poor response time. Tools, particularly editors, are inadequate because of the constraints on their design imposed by the mainframe environment. For example, character-at-a-time interaction with a mainframe editor is prohibitively inefficient, both because of the channel and interrupt structure of these machines, and because of the large number of users typically attached. The best compromise in such an environment is a full screen editor that relies on some local editing capability in the terminal, which in turn communicates with the mainframe a screen at a time. Furthermore, editing on a mainframe system, particularly one that is simultaneously running a production load, is often plagued by poor response time. In short, mainframe editors cannot approach the functionality and responsiveness of the best personal computer screen editors.

Current technology opens another possibility, namely integration of personal computer workstations with the mainframe system in such a way that program composition can be carried out on the workstation, while the mainframe provides a mechanism for managing and coordinating the work of several programmers. Many divisions of labor are possible. For example, compilation might be carried out by submitting a mainframe job, or on the workstation with a compatible compiler. Ultimately, of course, integration and testing have to be done on the target system, but even here the workstation could provide an easy-to-use mechanism for submission of jobs and retrieval of their output.

A particularly difficult problem is the development and maintenance of operating system software, because testing it usually requires stopping production and taking over the entire machine for "systems time," unless a dedicated development machine is available. However, the virtual machine concept that was discussed above can alleviate this problem. In such an environment, a new (version of the) operating system can be tested in its own

virtual machine, while production continues on another. At worst, a bug in the system under test can crash its own virtual machine; the physical machine and the virtual machine monitor are unaffected, and hence any other virtual machines will continue to run without interruption.

Chapter 8

Operating Systems

8.1 Introduction

In order to serve RLG's purposes, operating systems should be primarily oriented toward the management of on-line transactions; toward the management of very large data bases; and toward a distributed environment.

Unfortunately, most current systems were designed to be general purpose (i.e. both batch and on-line; both scientific and commercial, etc.); to provide time-sharing rather than transaction management services; to be oriented toward the processing of files rather than data bases; and to run in a centralized rather than a distributed environment. Transaction management, data base management, and distributed processing capabilities have generally been "grafted" on as afterthoughts, and as a result simplicity and integration have suffered. In addition, many systems were implemented prior to the recognition and acceptance of the software engineering principles discussed in the preceding chapter and are clumsy and difficult to maintain.

Some authors might have split this single chapter into four others: Operating Systems, Transaction Monitors, Data Base Management Systems, and Data Communications Systems, reflecting the structural separation in current system software offerings. We have not done so. While each component has an individual identity and character, we belive that each should be by design part of a well-integrated whole. Evidence exists that others share that view, and have done so for some time. [97] [98].

Unfortunately, no tangible example of such an exemplary system has yet reached the marketplace, notwithstanding the passage of time and an increasing consensus regarding the requirements and the need. To make matters worse, library information processing has some unique qualities which have placed it near the edge of the mainstream, causing in turn additional mismatches between

---

97. Gray, James N., "Notes on Data Base Operating Systems", Operating Systems: An Advanced Course, Bayer, R.M. et al. (Eds.) Springer-Verlag, Berlin Heidelberg New York 1979

98. Stonebraker, Michael, "Operating System Support for Data Base Management", Communications of the ACM, July, 1981, p 412

requirements and functions currently supported in off-the-shelf system software products.

The remainder of this chapter will briefly explore RLG's present and future operating system requirements; will survey some operating systems currently available; and will offer an assessment of what the future might hold.

8.2 Requirements of an Operating System

8.2.1 Transaction Management

The transaction management component must be capable of creating and scheduling transaction instances to act upon incoming messages. On-line transactions in library applications tend to be conversational, rather than unitary. This is particularly true in the on-line catalog application, where user-machine dialogs can be quite protracted. Thus, the system must be capable of maintaining user state across message boundaries. The system must also support batch transactions, and must allocate resources to all types of transactions in a manner appropriate to their response requirements. The overhead associated with transaction creation and scheduling and with task switching must be low.

The transaction management component must provide support for the maintenance of data base integrity in spite of abnormal events such as software errors, system failure, power outages, etc. Ways to guarantee the atomicity of data base updates associated with a transaction (Begin/End/Abort Transaction) must be supported. Transaction logging with efficient roll-forward recovery is an additional necessity.

Concurrency control is also a responsibility of the transaction manager. This includes the ability to set and modify multiple locks, with deadlocks detected by the system. The lock manager must be highly efficient, since locking is a common occurrence in systems of this kind. Locking in general-purpose operating systems is sometimes very costly.

8.2.2 Data Base Management

A prior part of this report has documented a hypothetical model for a local integrated library data base.[99] This picture implies the usual need for accomodating records, fields, and sets, with many inter-record and inter-set relationships.

---

99. "A Distributed System Architecture", Carnegie Study on Distributed Computing, Part II , page 17.

That portion of the model dealing with the administration of the library (i.e. orders, invoices, vendors, patrons, physical item control, etc.) would be admirably served by a data base management component which adheres to the relational data model. [100]

It unlikely that such a statement applies to heading descriptions, however, and even less to bibliographic descriptions While it is certainly true that the data can be accomodated by the relational model, the power and generality of the model is inappropriate in this case, and a severe performance penalty is paid. Moreover, most relational systems do not handle variable-length elements, which is a concomitant of bibliographic and heading data. The need to store and retrieve variable information will grow greater, given the movement toward full text and document-oriented applications.

Another stumbling block is a requirement to support efficient retrieval using boolean search expressions containing words derived from one or more different fields. Most existing data base systems fail to provide the requisite capabilities in this area as well.

8.2.3 Dictionary Management

The names and descriptions of the data bases, the files, the transactions, the network, and the users need to be stored somewhere by the operating system. Generally, this is accomplished ad hoc by the operating system using software particularized for each purpose. A far more elegant approach is to use the capabilities of the data base management component for this purpose. Much code is shared thereby, and a uniform interface is presented to all those interacting with the system. [101]

8.2.4 Distributed Processing

The foregoing discussion dealt with services to be provided by the operating system in a single host environment. In a distributed environment, those same services must be provided across the several hosts within a distributed processing network. This implies in the first instance an efficient and well-designed set of protocols to support the inter-host transfer of requests, system messages, and data. Modern computer networks are designed in a highly structured way. [102] In particular, the American National Standards Institute has specified the International Standards Organization's Open Systems Interconnection Reference Model as the structure to be adhered to in the U.S. Accordingly, RLG requires that the ISO model form the basis for any inter-system linkage in the RLIN network.

---

100. Codd, E. F., "A Relational Model of Data for Large Shared Data Banks", Communications of the ACM, Vol. 13, page 377.

101. Gray, James N., op. cit.

102. Tanenbaum, Andrew S., Computer Networks, Prentice-Hall, Inc., 1981

Notwithstanding the ability for two or more hosts to communicate, some of the problems dealt with in the single host environment become substantially more difficult in the distributed case. The provision for transaction atomicity in the case where data residing in several locations must be updated is a case in point. Another is the provision of locking facilities and deadlock detection across the network. Dictionary management must be extended to include location information; this must be universally accessable within the network.

Fortunately, the need within RLG for such concurrency and consistency control functions across host systems does not exist in the short and medium term, as long as simple information retrieval and batch record transfer are deemed sufficient. In the longer term, however, it will be necessary to have full support for these functions if a transition to workstation-based systems is to occur.

8.2.5 Other Requirements

If distributed system components are to be deployed when and where they are needed, then the system must not require more than a single staff member for each shift of operation. The installation of the system should take no more than a few hours, and its subsequent maintenance accomplished via data communications links between the system and the vendor.

The system should be hospitable to software developers. It should support a variety of higher-level languages (e.g. Pascal, Ada, "C"), including syntax-directed editors; a powerful debugging environment, and a non-interfering test environment to allow production and development to proceed concurrently.

8.3 Survey of Current Operating Systems

Stonebraker rightly concludes that operating system services in many existing systems are either too slow or are inappropriate for serious transaction and data base management. Current data base management systems usually provide their own services and make little or no use of those offered by the operating system. He remarks further that a small efficient operating system with only desired services is far more preferable, and that only the so-called "real time" systems currently fit that mold. [103] The following survey readily supports this point of view.

---

103. Stonebraker, Michael, op. cit.

8.3.1 UNIX

UNIX [104] [105] exists in many forms and versions, but is primarily licensed by the Bell Laboratories. UNIX is becoming a de facto standard among small machine operating systems. It is largely portable, and now runs on the largest IBM 370s, on many personal computers, and on most popular mid-range machines. What is behind its popularity? Its foremost attributes are simplicity, elegance, a superb development environment, and close adherence to software engineering principles. Unfortunately, it is singularly inappropriate for transaction and data base management. It lacks a file system that can underly a data base system, and does not provide functions that guarantee transaction atomicity. Data base recovery is likewise not provided for. UNIX also lacks a shared memory segment for state information such as locks, etc. At least one implementor has successfully used UNIX as a data base operating system, but at the expense of having to add many of those functions which are too inefficient or missing altogether. A commercial variant of UNIX called XENIX [106] provides a shared memory segment, crash recovery at the file level, and extended language support. [107]

8.3.2 VMS on the VAX

The VAX family was designed and developed by the Digital Equipment Corporation as a successor to the PDP-11. The design goals were oriented toward having a successful line of general-purpose computers with a bias toward timesharing, networking, and scientific computation. As such systems go, it, like UNIX, is extremely well-designed and implemented. With the exception of the RMS file system, it lacks many of the functions needed to support transaction and data base management. Furthermore, the VAX lacks the requisite capacity to perform large amounts of high-speed data transfer.

8.3.3 MVS on 370-Compatible Machines

MVS is decended from a long line of batch-oriented operating systems. By itself, it is lacking most necessary functions to support transaction and data base management, and what does exist is painfully slow and expensive. MVS requires at least one dedicated system programmer to maintain it, and it is difficult to operate.

---

104. UNIX is a trademark of Bell Laboratories

105. Ritchie, D. M., and K. Thompson, "The UNIX Time-Sharing System," Bell System Technical Journal, Vol. 57, No. 6, Part 2, July-August 1978, page 1905.

106. XENIX is a trademark of Microsoft, Inc.

107. Deitel, Harvey M., Operating Systems, Addison Wesley, 1984

Two add-on components to MVS exist, either of which may be used to support on-line commercial data base applications: IMS/370 and CICS. The former supports transaction and data base management for complex hierarchies of fixed-length segments. This model is not appropriate for bibliographic data.

CICS provides a very limited set of transaction management functions (scheduling) but no functions to assure data base consistency and recoverability. Conversational transactions are supported, but with a large increment of CPU overhead. The VSAM file system provides a base (barely adequate) on which the record management component of a data base management system could be built. Some library information systems have made limited and careful use of third-party transaction and data base management components in conjunction with MVS to compensate for the inadequacy of the IBM offerings. [108]

8.3.4 VM on 370-Compatible Machines

VM is not an operating system, if one applies a strict definition of the term. Rather, it is a monitor whose function it is to implement two or more error-free virtual machines on top of a single real machine. Operating systems capable of execution on 370-compatible machines are then run in the virtual machines (see above). A single-user system, the Cambridge Monitor System (CMS) is often used in a virtual machine to provide a user interface to a typical set of objects (files, compilers, an editor, etc.) CMS supports a version of VSAM, but otherwise is lacking in the other required capabilities.

8.4 An Assessment of the Future

The bleak assesment in this chapter is in marked contrast to the advances in function and performance noted in previous chapters dealing with hardware technology. It is also at odds with the growing recognition and acceptance of the software engineering principles discussed in the previous chapter. This section offers an opinion on what can be expected over the next several years in this crucial arena.

8.4.1 Research and Development Directions

Cynics have observed that the issues of concurrency control and recovery in a distributed environment have provided sufficiently fertile ground for publication that more than a few of today's computer science faculty owe their tenure to the phenomenon. Likewise, the flow of observations on and extensions to the relational data model seem endless. This will doubtless

---

108. "Evaluation of Local System Alternatives", Carnegie Study on Distributed Computing Part IV, page 15.

continue, since relational systems constitute the major direction for commercial data bases of the future.

Information (i.e. text) as opposed to the formatted data in a corporate data base has been ignored for the past decade. The papers published on information retrieval during this period can be contained in a very thin binder. A change is in the wind, however. This is caused by the recognition of office automation and electronic publishing as significant future markets. The kinds of information to be dealt with in these applications is likely to give rise to data base software capable of handling both operational and bibliographic library data.

8.4.2 Vendor and User Reluctance

An enormous number of commercial applications are now in existence, representing a multi-billion dollar investment by their implementors. Most of these are not portable, either because they pre-date the existence of software engineering as a discipline, or because the principle of portability is not deemed important in the commercial data processing world.

An understandable reluctance exists on the part of the larger hardware vendors to introduce new operating systems which would render this body of applications obsolete, and force a prohibitively expensive conversion to the new technology. Furthermore, users of the present systems do not in general perceive a need to change in spite of their high cost and complexity. It is considered to be simply a cost of doing business. The market for an operating system that is ideal for RLG's needs has yet to be developed.

One highly respected market surveyer has observed that the time is ripe for such a product to be brought forward, and that it would doubtless find acceptance within five years. [109] Until this occurs, progress in the automation of libraries otherwise made possible by gains in hardware function and price/performence will be hampered.

---

109. "Amdahl Corporation's ASPEN: an Idea Whose Time Has Come", presentation by the Gartner Group, Inc., Stamford, CT.

Carnegie Study on Distributed Computing

PROCESSING AND DATA DISTRIBUTION WITHIN THE RESEARCH LIBRARIES INFORMATION NETWORK

PART VI

THE FINANCIAL IMPACT OF DISTRIBUTION

November 1983

The Research Libraries Group, Inc.

Stanford, CA 94305

Table of Contents

Chapter 1 Introduction 1

Chapter 2 A Financial Model of the Data Resource Node 3

2.1 Modeling Tools 3
2.2 Scope of the Model 4
2.3 Sources of Data 4
2.4 Structure of the Model 5
2.5 Choosing the Initial Scenarios 5
2.6 Invariant Assumptions 7
2.7 Scenario 1 - Inertial 11

2.7.1 Summary of Assumptions 11
2.7.2 Machine Resource Forecast 11
2.7.3 Income/Expense Forecast 11

2.8 Scenario 2 - Conservative Change 15

2.8.1 Summary of Assumptions 15
2.8.2 Machine Resource Forecast 15
2.8.3 Income/Expense Forecast 15

2.9 Scenario 3 - Aggressive Change 19

2.9.1 Summary of Assumptions 19
2.9.2 Machine Resource Forecast 19
2.9.3 Income/Expense Forecast 19

2.10 Initial Conclusions 23
2.11 A Fourth Scenario - Conservative Prime 23

Chapter 3 Institutional Library Automation Costs 27

3.1 Comparative Local System Costs 27
3.2 Total Institutional Costs 29
3.3 Conclusions 29

Appendix A Detailed Scenario Assumptions 33

Chapter 1

Introduction

Assessing the financial impact of distribution is one of the most important goals of the Carnegie Study. In other areas of the project, staff have gathered or derived much information of use in projecting an initial view of RLG's financial future given a distributed systems environment. Examples are:

- The architecture for a distributed RLIN network
- The general requirements for a local system
- Member plans for
    * Installing local processing capabilities
    * Retrospective Conversion
    * Increasing or decreasing system activity levels
- The existence of local system alternatives that are compatible with the architecture, and in particular
    * Their functionality
    * Their cost
- An Assesment of Future Technology which focuses on
    * Different ways of looking at things
    * Availability of new technology
    * Changes in price/performance.

No financial assessment can be considered adequate unless the kinds of information cited above are carefully taken into account. The financial modelling effort drew heavily upon this body of information.

Chapter 2 describes a financial model that was formulated to study the effects of distribution upon the operational finances of the data resource node (that is, the RLIN central site). The scope of the model is defined, as are the sources of data. The rationale for choosing the initial scenarios, and the assumptions held in common among the scenarios are discussed. The results derived from three model runs using different scenarios are described and contrasted.

Chapter 3 deals with representative costs likely to be incurred by a member institution if processing occurs locally. The information in this chapter is largely derived from pro forma system configurations specified by vendors in the RFI process.

Chapter 2

A Financial Model of the Data Resource Node

2.1 Modeling Tools

The term "model" when applied to the design of airfoils or to weather prediction implies the use of complex, sophisticated mathematics to abstractly represent physical phenomena over time.

Models that are constructed to predict financial outcomes are complex also, but not usually in a mathematical sense. The mathematics are mostly simple arithmetic, with occasional use of elementary analysis techniques. The complexity arises from the number of pieces that can vary in different ways, and particularly from the frequent lack of "hard" information to use as input to the model.

Once the model has been constructed and initial data have been gathered or invented (artistically and intelligently, one hopes) the model can quickly perform the thousands of simple arithmetic operations necessary to produce the "financial picture" over a period of time. More importantly, the user of the model should be able to change the input data to see how the picture changes. Different sets of input data that produce different pictures are called scenarios.

For the problem at hand, the project staff had to make a choice of implementation tools. The alternatives were reduced to three:

1. A Procedural Language: Project staff would implement the model using a language such as BASIC or APL. The former could be used on a personal computer with appropriate memory and secondary storage; the latter could be used on the RLIN IBM 3081.

2. EFPM: This is essentially a generalized model, which when supplied with the appropriate parameters and data, can be used to model a large class of problems. Most applications of EFPM have been oriented toward large institutions. The system is offered as a time-shared service at Cornell University's Computation Center. EDUCOM has played a strong role in popularizing its use.

3. A Spreadsheet Package: Spreadsheet packages have sprung up over the past two years, and are among the most popular software products available for personal computers. They allow arrays of numbers to be conveniently input, manipulated using complex arithmetic expressions, stored, and retrieved. The packages usually provide an integrated graphics capability to simplify the representation of results.

The first alternative was discarded because of the need for a trained programmer. We wanted our financial analyst to be able to interact directly with the model. The second alternative, although considered appropriate for the task at hand, was rejected because of logistical problems in the use of a remote computing resource. The third was chosen because of ease of use, the need to have project staff able to work at odd hours and at home, and because of a good functional fit between the tool and the model to be constructed.

The particular package chosen was Lotus 1-2-3, which was designed to run on the IBM PC. The staff's experience with 1-2-3 as a vehicle for implementing this model will be of interest to some, but is beyond the scope of this report. It will be the subject of a separate paper.

2.2 Scope of the Model

The model addresses operational income and expense. It addresses neither budgets associated with program and management support, nor grant-funded development activities. For the present, the model does not encompass data communications either, since the provision of RLIN telecommunications services has always been a separately budgeted activity "sitting on its own bottom".

2.3 Sources of Data

The basis for the model is a set of projections from the current state: static system activity except for that implied in the membership growth projections; costs inflated using the Stanford Long-Range Forecast (SLRF); and pricing structure similarly treated.

Applied to this base are those changes implied from information gathered in the Member Survey: planning dates for movement of cataloging and acquisition processing from the central site to local systems; and estimated retrospective conversion activity over time (See Part III). Additional information from central staff estimates is used in instances where information was not sufficiently firm to permit meaningful survey questions to be framed.

Production statistics provide important input: computer resource measurements taken from the current version of RLIN allow system capacity requirements to be computed on alternative workload mixes.

The Long-Range Technology Assesment (See Part V) provides projections of configuration price/performance that allow hardware replacement strategies to be assumed, together with their concomitant expense.

2.4 Structure of the Model

The model can be most simply characterized as a network of spreadsheets. Figure 1 is a simplified schematic of the network. The two primary entities are the EXPENSE MODEL and the INCOME MODEL; each is fed data successively from several other nodes of the network. The primary sources of data are described above.

2.5 Choosing the Initial Scenarios

Choosing initial scenarios is not easy given the large number of possibilities. The choices should reflect the goals of the organization; they should produce different results; and the results should be sufficiently interesting to provoke more questions, more scenarios, and more model runs. Our rationale for choosing the initial scenarios reported herein is described below.

The sought-after scenario is one which optimizes between two variables: degree of compensation for income lost because of the migration of processing; and the degree to which resource sharing is accelerated.

Compensation for lost income can be realized through one or some combination of four management actions:

1. Cost reduction, or
2. Price increase, or
3. Volume increase, or
4. Offering services that are different in form, substance, or audience.

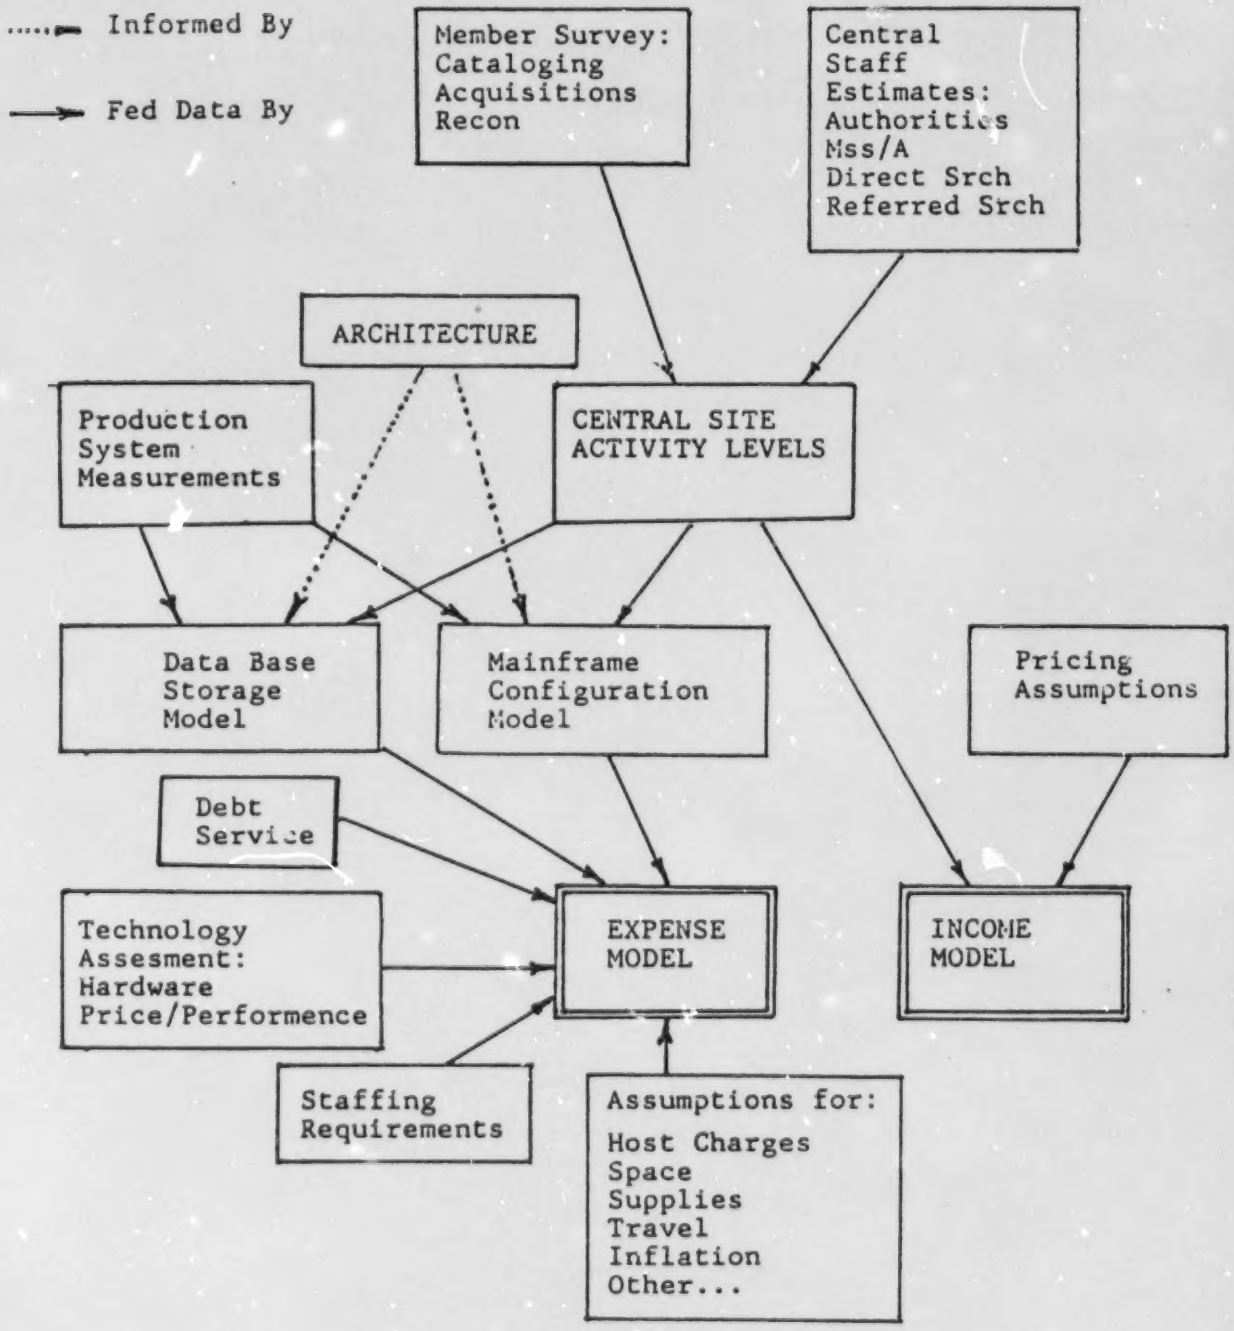

Figure 1.
Central Site Financial Model

The first is asserted to have limited effect, given the magnitude of fixed costs in computer hardware (price/performance improvements notwithstanding), and the critical mass of staff required to continue to offer central site services. The second is onerous for obvious reasons. The third seems remote, given the acquisition and staffing patterns likely to obtain in member libraries over the next decade. The fourth, by elimination, is worthy of initial investigation, particularly if the new services facilitate the further sharing of resources.

The new service that seems obvious for initial consideration is referred public service searching. This is defined to be search requests, issued either by reference personnel or by patrons, which have not produced "hits" locally and are therefore referred to the data resource node. It can be expected that such a service would result in a significant increase in interlibrary loan requests, which is the primary quantitative measure of resource sharing activity. The three scenarios presented below differ primarily in the degree of agressiveness displayed in providing capacity and support for this new service. Parenthetically, the "new" service is really a re-packaged service that already exists, targeted to a much wider audience.

2.6 Invariant Assumptions

The primary assumptions that remain unchanged across these three scenarios are as follows:

1. The central site activity levels implied in the Member Survey, i.e. cataloging, acquisitions, and retrospective conversion.

2. Given 1. above, the growth pattern of the data base. See Figures 2,3, and 4 for representations of data base growth between fiscal years 1984 and 1992.

3. Resource consumption assumptions taken from production system measurements.

4. Pricing of existing services. New services are pegged to existing rates to the extent possible.

5. The existing membership growth scenario.

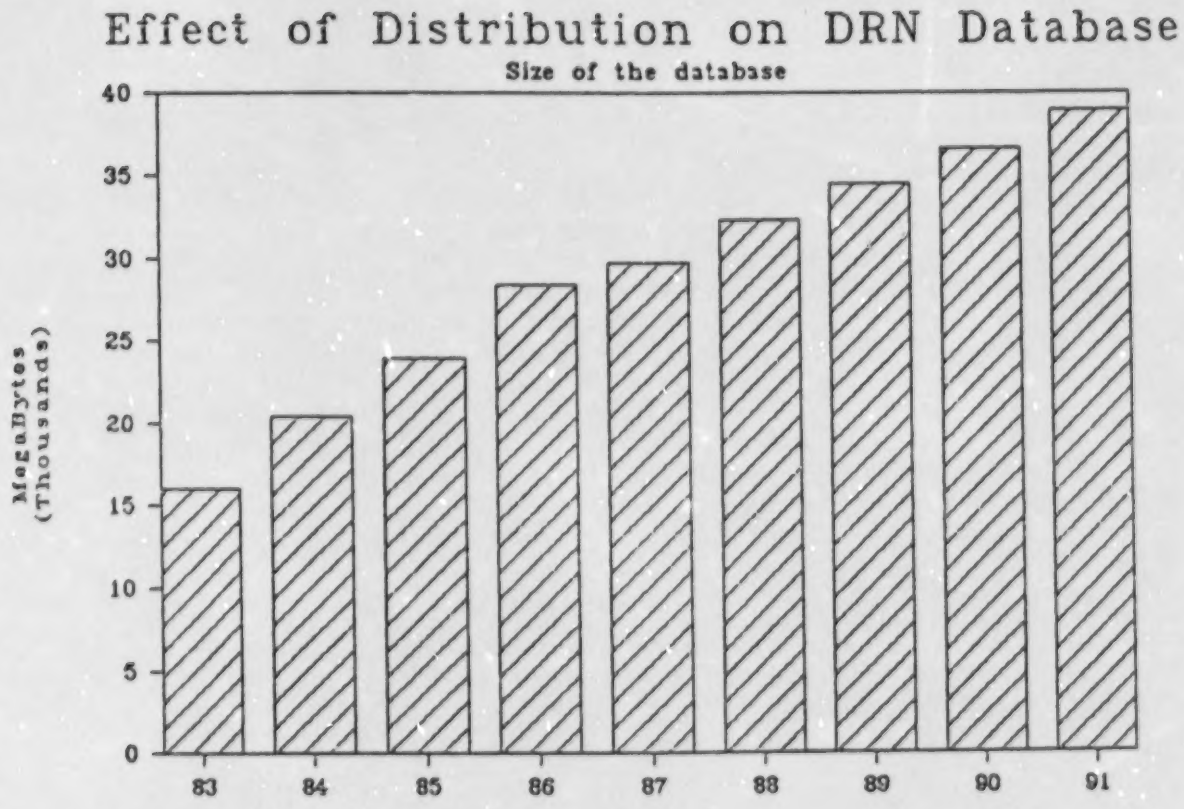
Effect of Distribution on DRN Database
Size of the database
MegaBytes (Thousands)
40
35
30
25
20
15
10
5
0
83
84
85
86
87
88
89
90
91

Figure 2.

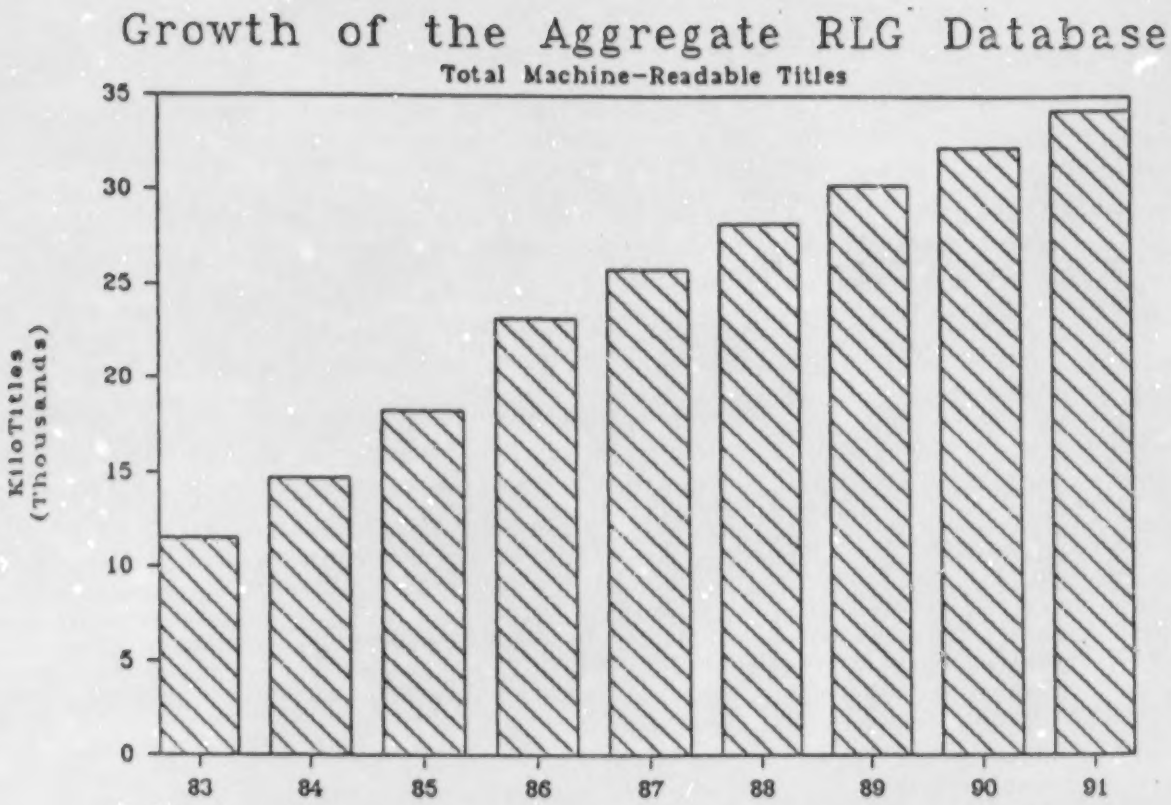
Growth of the Aggregate RLG Database
Total Machine-Readable Titles
KiloTitles (Thousands)
35
30
25
20
15
10
5
0
83
84
85
86
87
88
89
90
91

Figure 3.

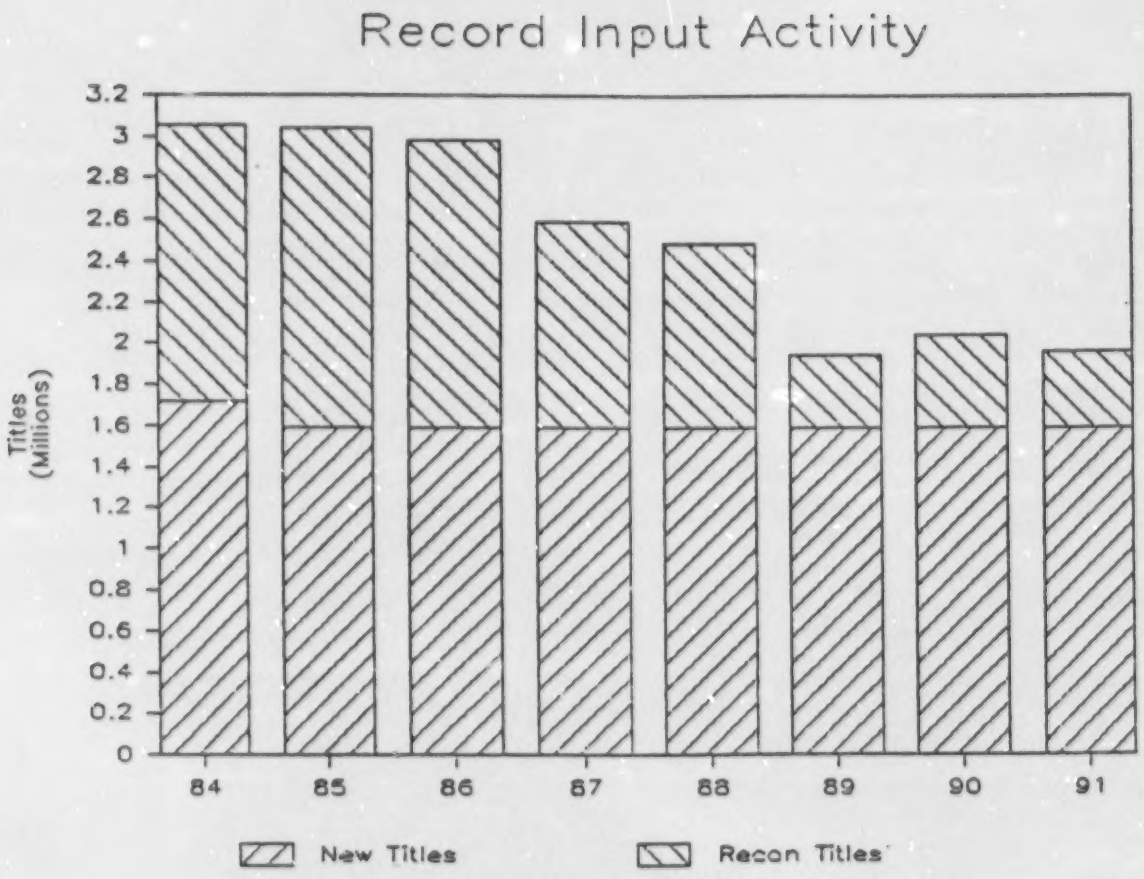
Record Input Activity
Titles (Millions)
3.2
3
2.8
2.6
2.4
2.2
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
84
85
86
87
88
89
90
91
New Titles
Recon Titles

Figure 4.

2.7 Scenario 1 - Inertial

2.7.1 Summary of Assumptions

This is the "let nature take its course" scenario. The central site configuration undergoes periodic modification to accomodate projected increases in data base storage requirements. Referred public services searching is supported to the capacity limits of the current CPU. Other changes reflect inflation as projected in the SLRF. Detailed assumptions for all three scenarios are contained in Appendix A.

2.7.2 Machine Resource Forecast

Figure 5 shows the CPU hours consumed within the limits of the present IBM 3081 Model K. The present version of ORVYL will consume up to one-half the machine until June, 1985, when the multiprocessor version of ORVYL will be ready for production. At that time, a 60% increase in CPU capacity will be available. The heavy black line shows the approximate capacity of the machine in CPU hours per quarter. This line is placed at the level of 60% of the hours available. A significant shortage in CPU resources is projected for the first six months of calendar 1985 if reference searching is allowed at projected demand.

Figure 6 shows the proportion of machine resources going to each class of service: Cataloging, Acquisitions, Authorities, RECON, Referred Technical Processing Searching (Rtp), and Referred Public Service Searching (Rps). Even given the restrictions of the 3081K, it is seen that the largest increment of capacity will be used for Rps searching.

2.7.3 Income/Expense Forecast

Figure 7 shows income and expense for fiscal years 1984-1991. It is seen that significant budget deficits will exist each year of that period. The horizontal lines drawn in the expense bars show that portion of expense attributable to repayment of long-term debt principal, and that this occurs in years 84-86 and 91. It can be seen that a deficit would exist each year even without the need for debt repayment. The cumulative deficit for the period is $14.5 million.

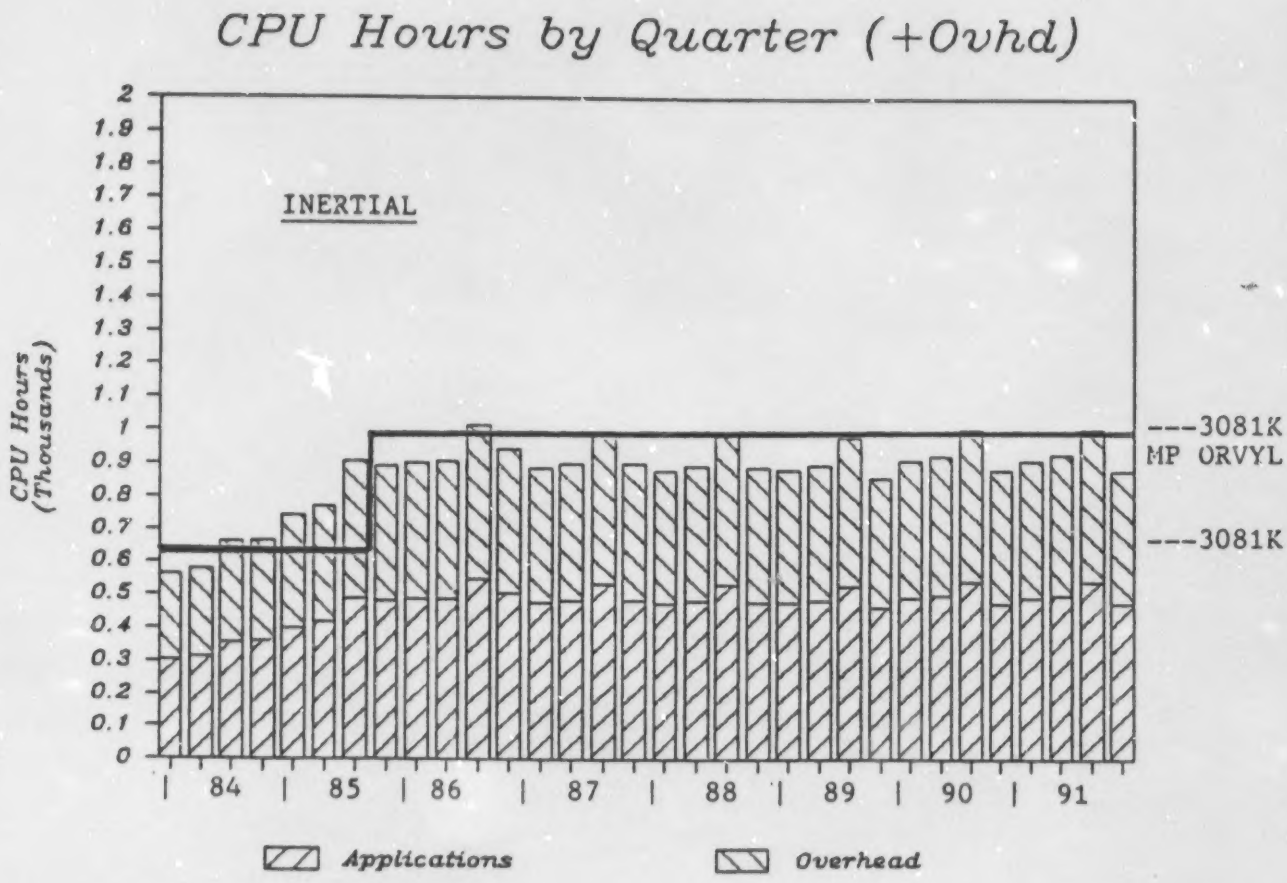
CPU Hours by Quarter (+Ovhd)
INERTIAL
CPU Hours (Thousands)
2
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
84
85
86
87
88
89
90
91
---3081K MP ORVYL
---3081K
Applications
Overhead

Figure 5.

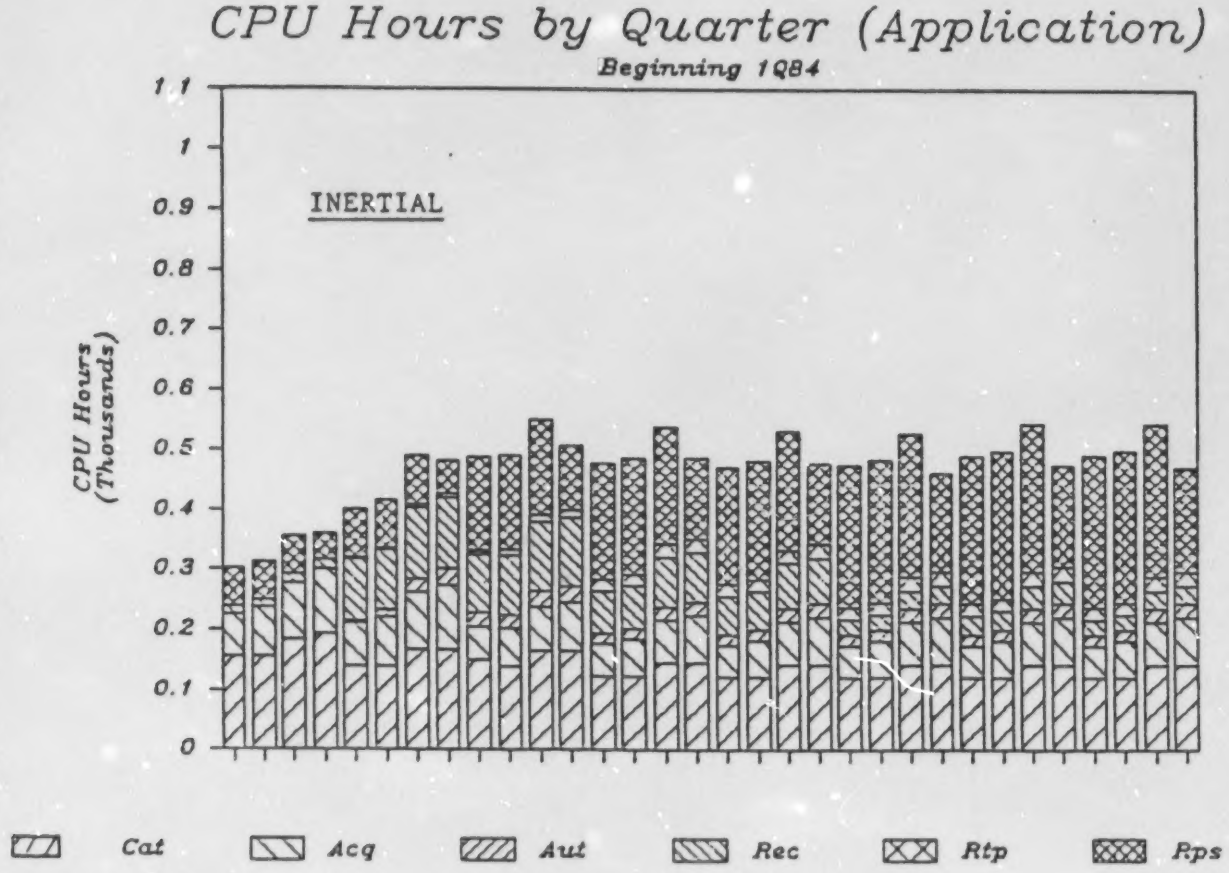
CPU Hours by Quarter (Application)
Beginning 1Q84
INERTIAL
CPU Hours (Thousands)
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Cat
Acq
Aut
Rec
Rtp
Rps

Figure 6.

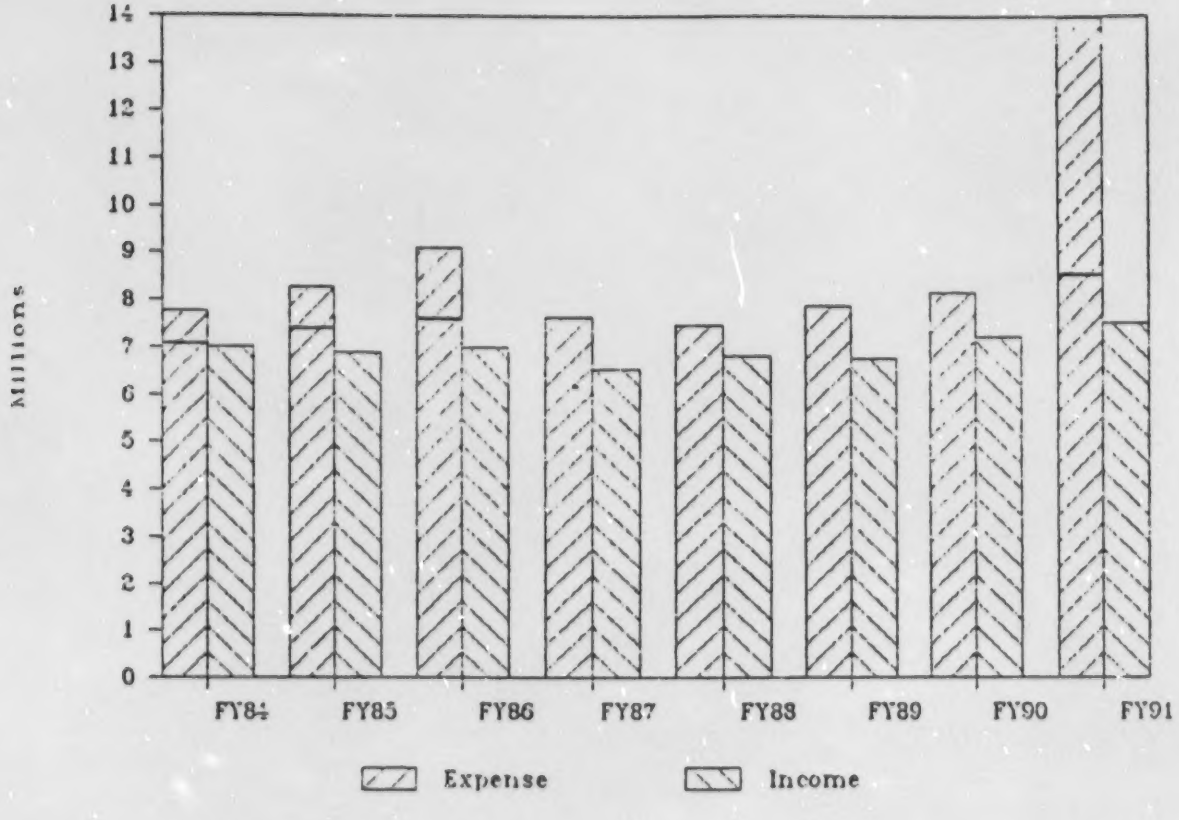
Distributed Processing Financial Model
Inertial Scenario
Millions
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
FY84
FY85
FY86
FY87
FY88
FY89
FY90
FY91
Expense
Income

Figure 7.

2.8 Scenario 2 - Conservative Change

2.8.1 Summary of Assumptions

This scenario is resonably conservative in what it attempts -- yet it succeeds in supporting an increase in referred public service searching. An early but modest upgrade in CPU capacity is assumed, together with those staff reductions considered possible given the need to maintain centralized services. Detailed assumptions can be found in Appendix A.

2.8.2 Machine Resource Forecast

Figure 8 represents the load versus capacity given this scenario. An upgrade to a 14 MIP uniprocessor is made at the beginning of FY 85. Aside from some tight 3rd quarters, the assumed load is relatively constant throughout the period. Figure 9 shows the proportional use of the system by each application. It is seen that referred public service searching is consuming more resources than the rest of the applications combined during the latter half of the period.

2.8.3 Income/Expense Forecast

Figure 10 represents forecasted income and expense for the eight year period. The horizontal lines drawn in the expense bars portray the amount attributable to principal payments to retire long-term debt. It is seen that except for 1984, income exceeds expense without debt service in the picture. The cumulative deficit for the period is $3 million.

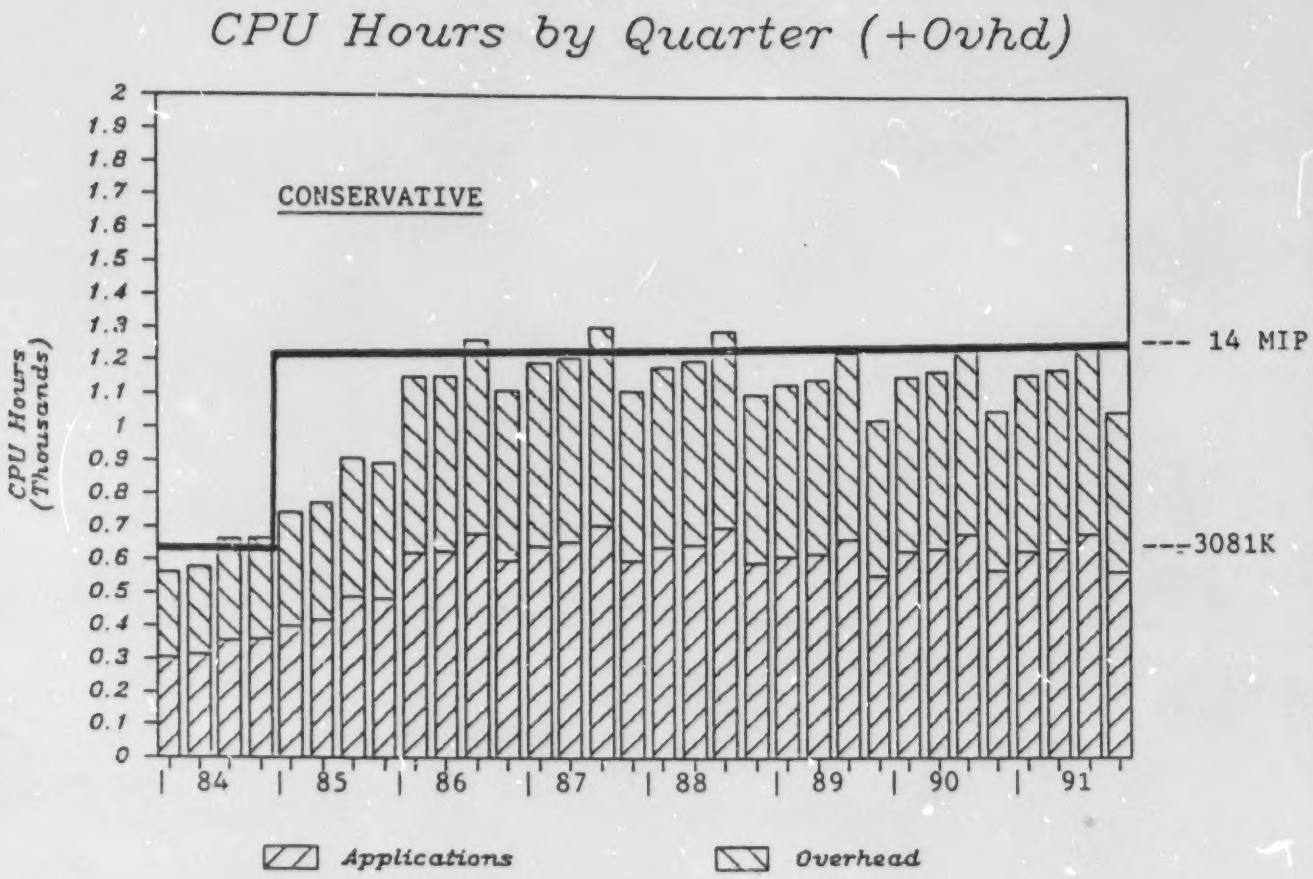
CPU Hours by Quarter (+Ovhd)
CONSERVATIVE
CPU Hours (Thousands)
2
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
84
85
86
87
88
89
90
91
--- 14 MIP
---3081K
Applications
Overhead

Figure 8.

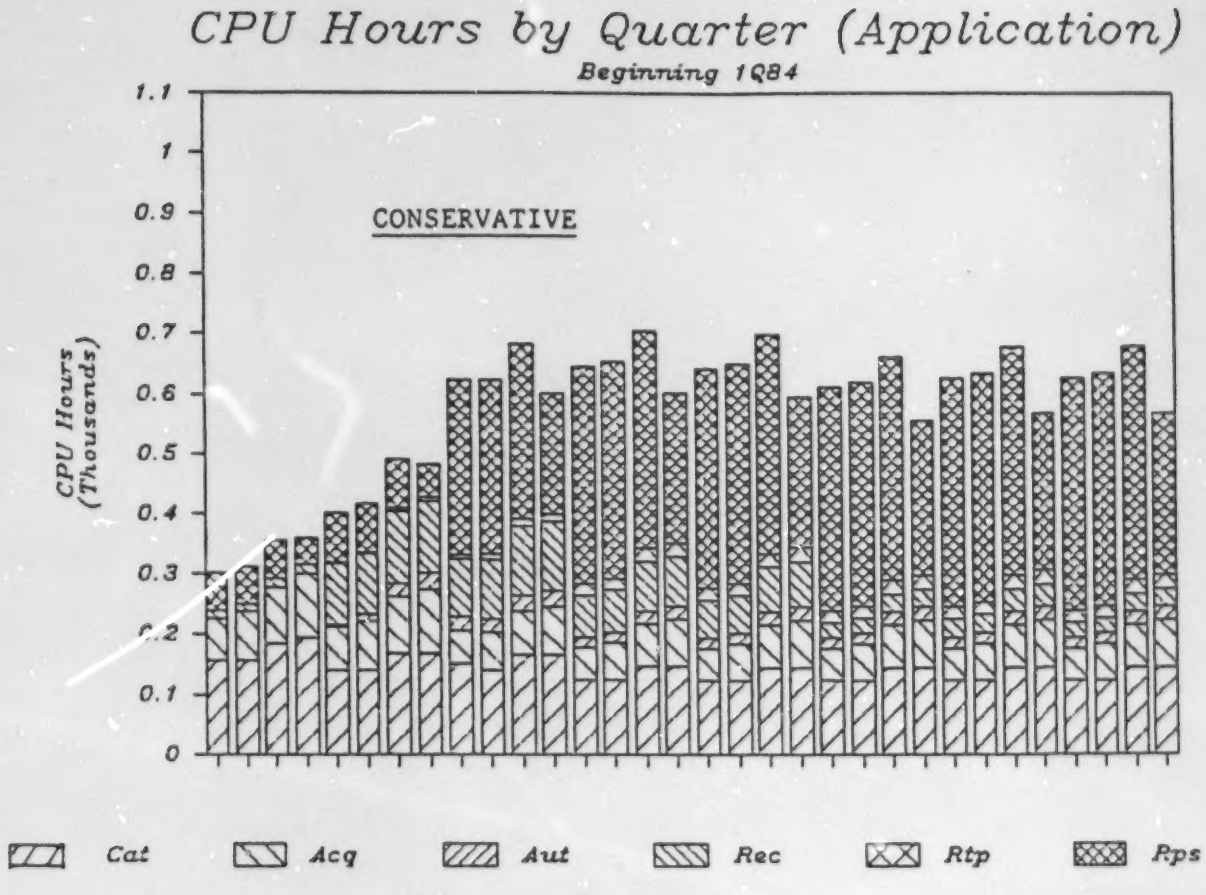
CPU Hours by Quarter (Application)
Beginning 1Q84
CONSERVATIVE
CPU Hours (Thousands)
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Cat
Acq
Aut
Rec
Rtp
Rps

Figure 9.

Distributed Processing Financial Model

Conservative Scenario

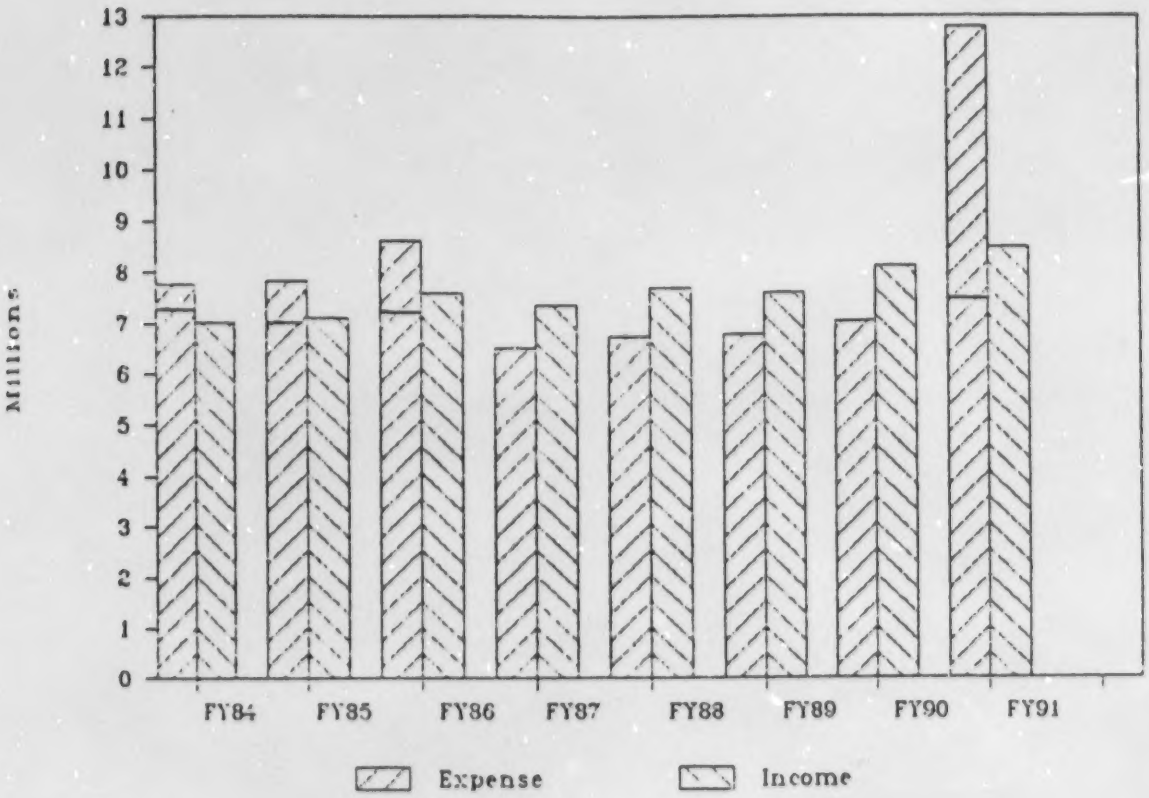

Figure 10.

2.9 Scenario 3 - Aggressive Change

2.9.1 Summary of Assumptions

This scenario aggressively adds system capacity, and assumes referred public service searching use to use most of the capacity. No staff reductions are contemplated in this scenario. Appendix A may be examined for the detailed assumptions.

2.9.2 Machine Resource Forecast

Figure 11 shows projected CPU utilization against capacity for this scenario. An upgrade to a 14 MIP uniprocessor is assumed at the beginning of FY 85; a further upgrade to 24 MIPs is assumed two years later. Note that the alternative path of staying on the 3081K with MP ORVYL is shown, with the need to reduce the assumed load in FY 86. Figure 12 shows the majority of the system resources going to support referred public service searching.

2.9.3 Income/Expense Forecast

Figure 13 shows income versus expense for the period. As before, the horizontal lines delineate the portion of expense devoted to principal payments. The cumulative deficit for the period is $1.2 million.

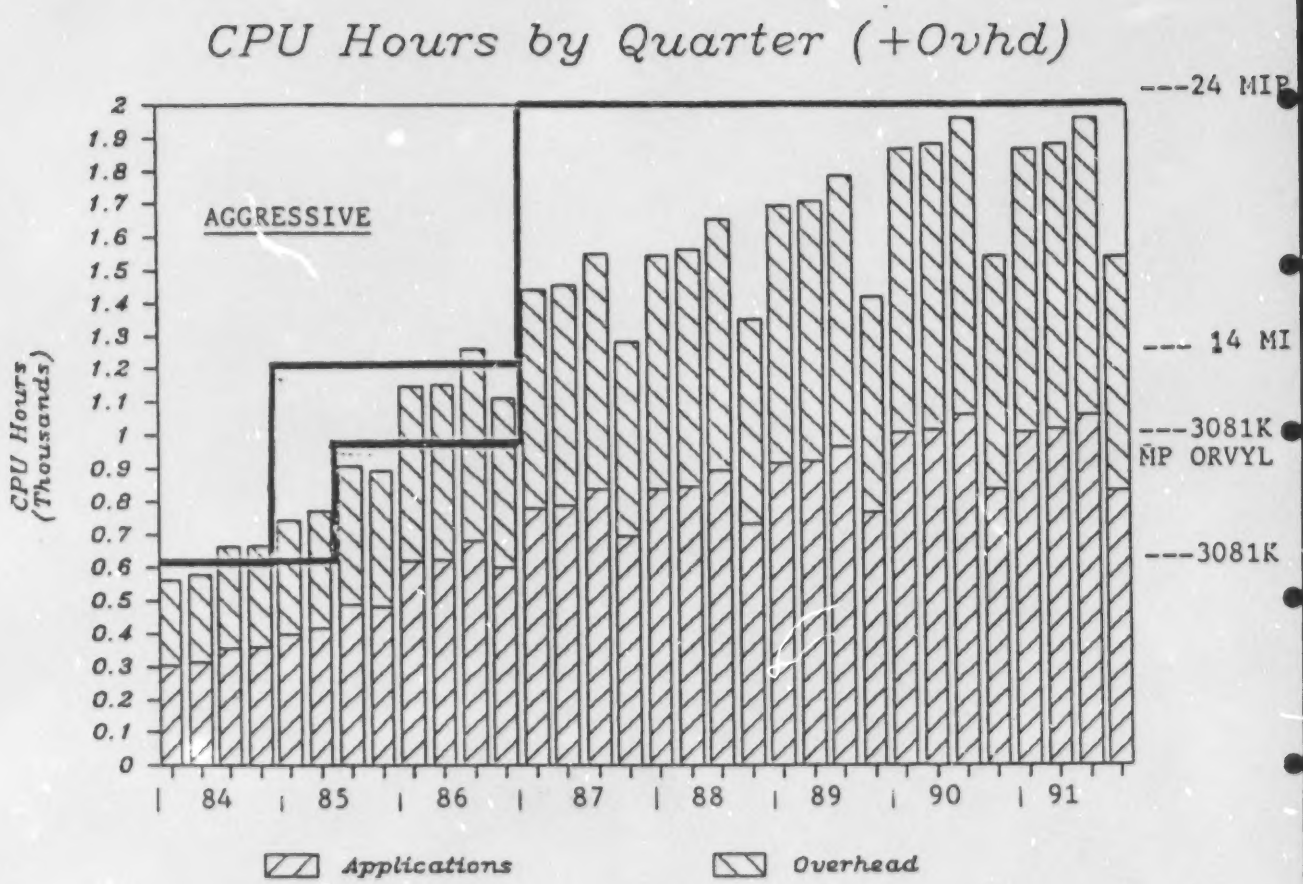
CPU Hours by Quarter (+Ovhd)
CPU Hours (Thousands)
AGGRESSIVE
---24 MIP
--- 14 MI
---3081K MP ORVYL
---3081K
2
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
84
85
86
87
88
89
90
91
Applications
Overhead

Figure 11.

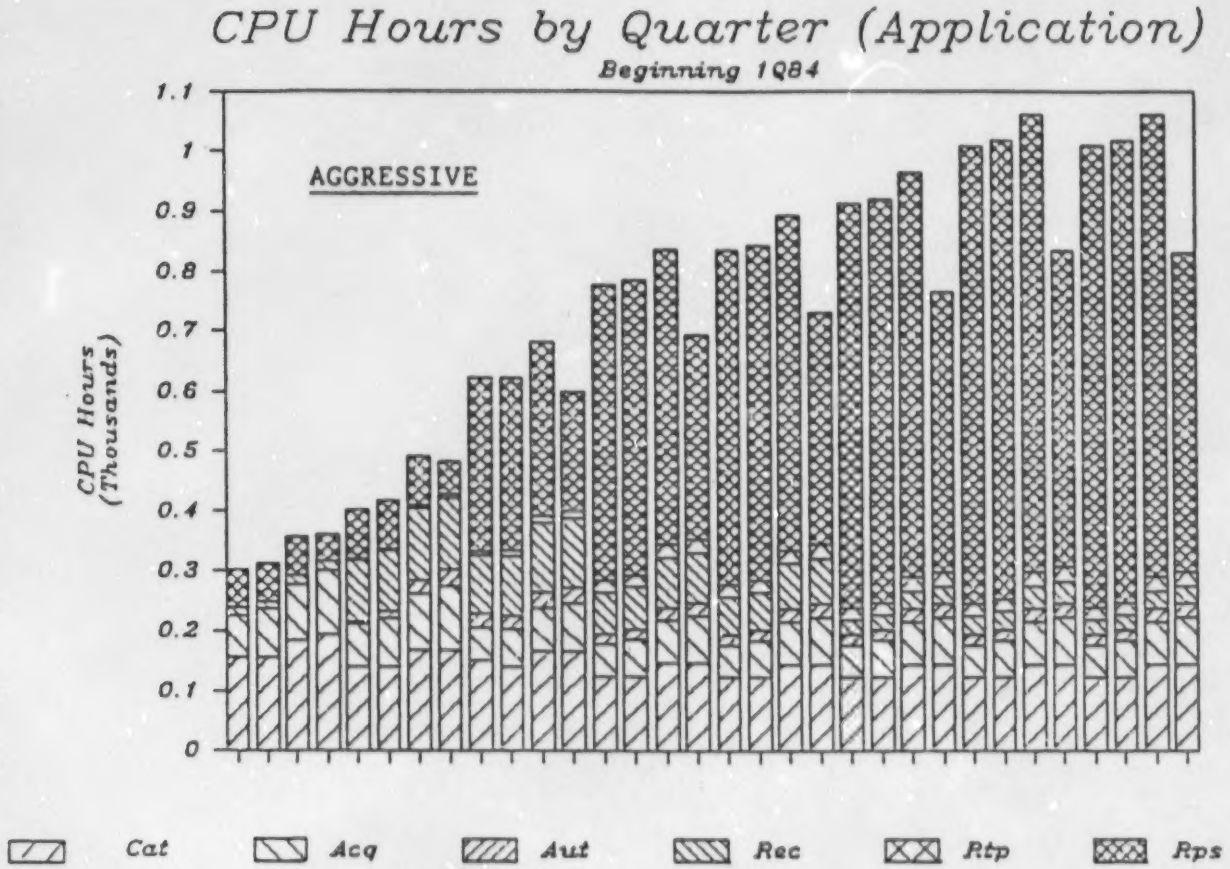
CPU Hours by Quarter (Application)
Beginning 1Q84
AGGRESSIVE
CPU Hours (Thousands)
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Cat
Acq
Aut
Rec
Rtp
Rps

Figure 12.

Distributed Processing Financial Model

Aggressive Scenario

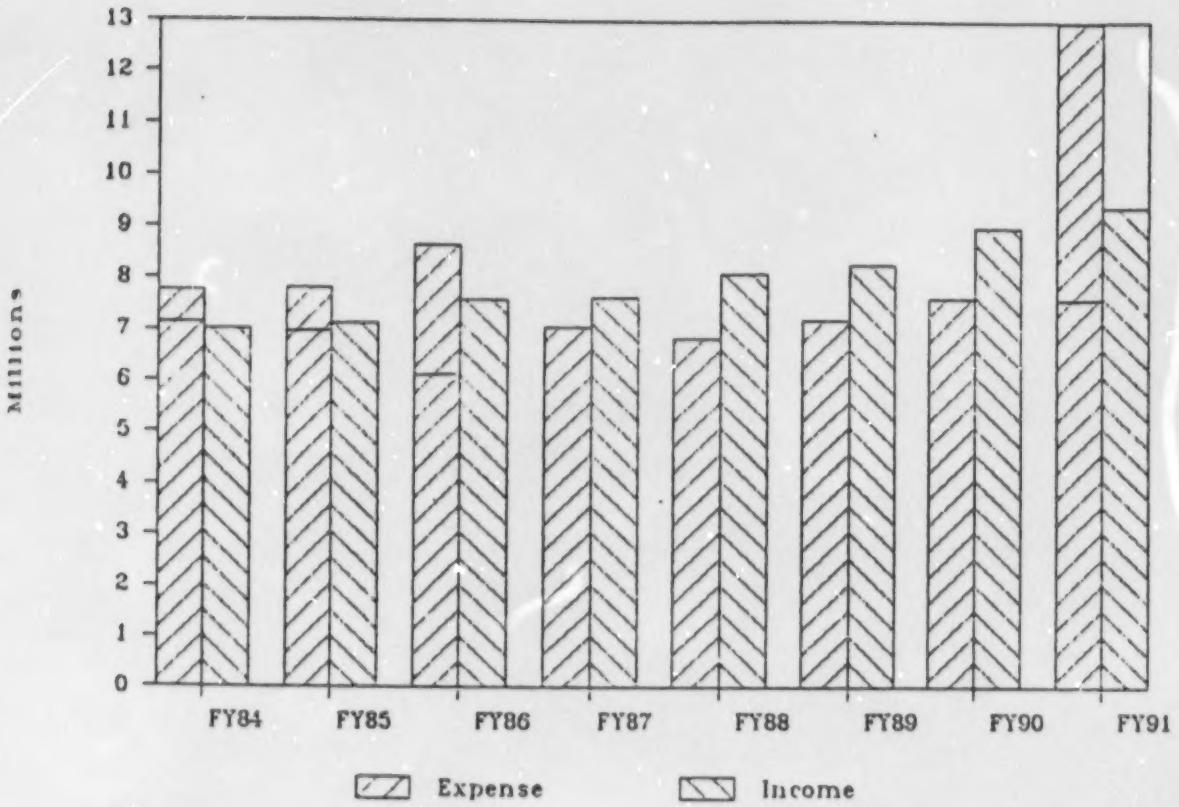

Figure 13.

2.10 Initial Conclusions

The Inertial scenario is seen to have a large cumulative deficit, and tells us what we have always known: that our long-term debt is a problem; that we need a larger base over which to spread our fixed costs; and that the loss of central site income due to distribution will further swell the deficit if no compensitory action is taken.

The Convervative scenario significantly reduces the cumulative deficit, but with some risk. The principal elements of risk relate to the speed with which local/DRN connections are established, and the degree to which spending levels by member institutions for referred searching can be supported. This may in turn suggest mechanisms for passing some costs through to the end user.

The Aggressive scenario further erodes the cumulative deficit, but the risks described above are significantly magnified.

Clearly, none of the three initial scenarios provide a satisfactory basis upon which to proceed, either because of deficits, risk factors, or both. We therefore re-examine our initial assumptions, particularly the feasibility of maintaining current pricing levels.

2.11 A Fourth Scenario - Conservative Prime

One pricing assumption that stands out as a candidate for change is the one for searching. The current rate, $0.13/CPU second, reflects the approximate marginal cost of providing the service, and has been used to encourage public service use of the system. It is viewed as a "bargain rate" which may be an unaffordable luxury in the future.

Other soft spots in the initial scenarios were the rapid onset of referred public service searching, and a very ambitious level of retrospective conversion in fiscal 1984. At the same time, the rate of movement of members to the use of local systems for cataloging and acquisitions seemed somewhat unrealistic, especially since intersystem linkage between the DRN and local systems will become available for limited use no earlier than the last quarter of 1985.

To address these observations, the income assumptions for the conservative scenario were modified to provide a fourth scenario: Conservative Prime. In this scenario, the rate for referred searching is $0.26/CPU second; referred searching levels for fiscal 1986 are significantly lowered; cataloging volume is increased by 200,000 per annum; and the erosion of acquisitions is delayed one year. The retrospective conversion estimates are moved out one year also.

Figure 14 shows income versus expense for the period. As can be seen, these modifications produce a cumulative budget surplus of $4.9 million over the eight-year period.

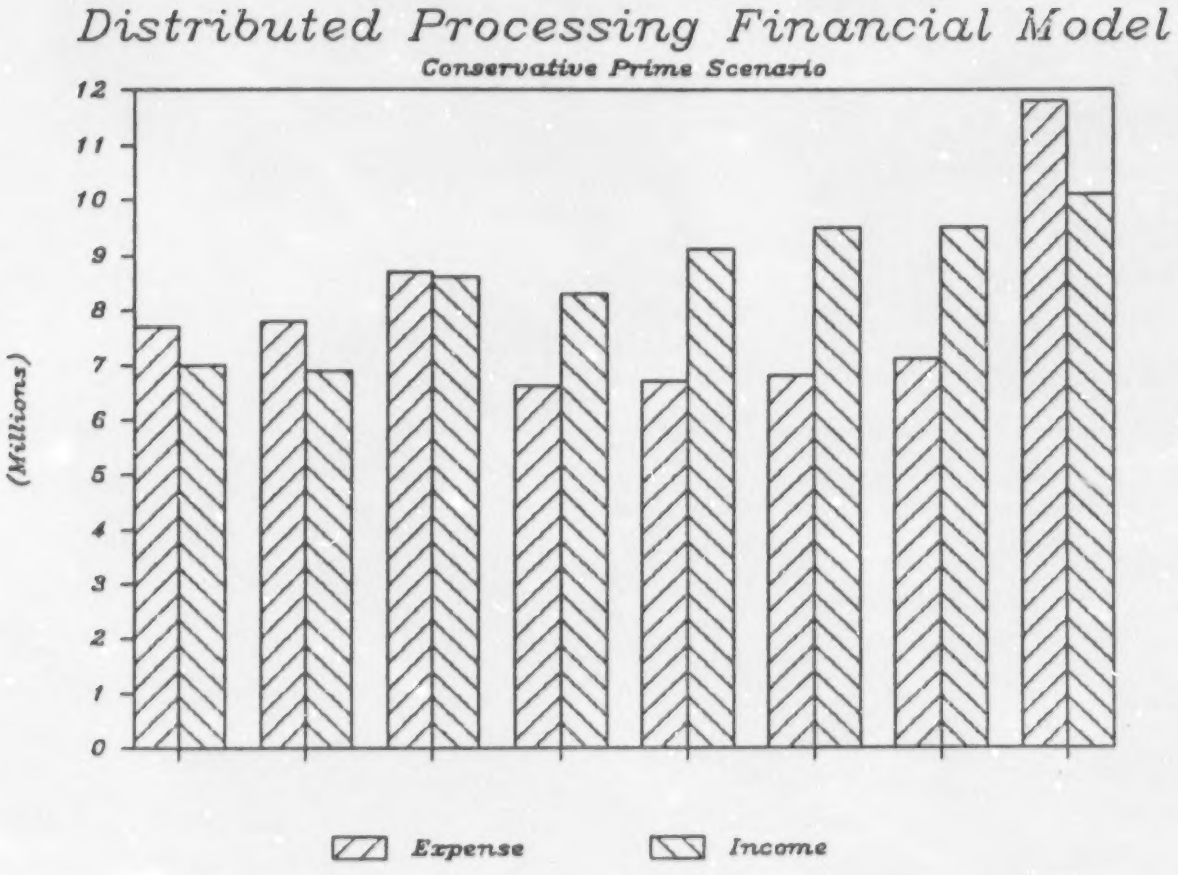
Distributed Processing Financial Model
Conservative Prime Scenario
(Millions)
12
11
10
9
8
7
6
5
4
3
2
1
0
Expense
Income

Figure 14.

Chapter 3

Institutional Library Automation Costs

The preceding assumptions about DRN utilization, both with respect to quantity and kind, together with estimated local system costs gained from the RFI responses, allow a rough model of local costs for an "average" RLG institution to be constructed.

It should be emphasized that this model is for illustrative purposes only. It is incumbent upon each RLG member institution to undertake a similar but more detailed analysis using assumptions that fit the particular case.

3.1 Comparative Local System Costs

One-time costs associated with the installation of a local system in this brief comparison include ten percent of the purchase price of the data center hardware for the local library system; fees to the system vendor for installation and training, and initial fees for software licensing. In some cases, such items as data conversion costs and custom programming will need to be included also. The costs associated with terminals, lines, and other communication equipment are likewise not included in this comparison, since these vary widely depending upon local geography. A reasonable approximation that may be used in the absence of a careful estimate of the "network" hardware component is an amount equal to the data center hardware component. Figure 15 compares one-time costs for the three recommended systems, assuming a configuration for an average library [1].

Annual recurring costs to support the local system include amortization of the remaining data center hardware purchase cost; staff salaries and benefits; software maintenance; and data center hardware maintenance. As explained above, the network costs (terminals, lines, modems, etc.) are not included, but data center amortization and maintenance costs provide a reasonable approximation of maintenance on the terminals and other equipment.

---

1. Carnegie Study on Distributed Computing, Part IV, p. 66

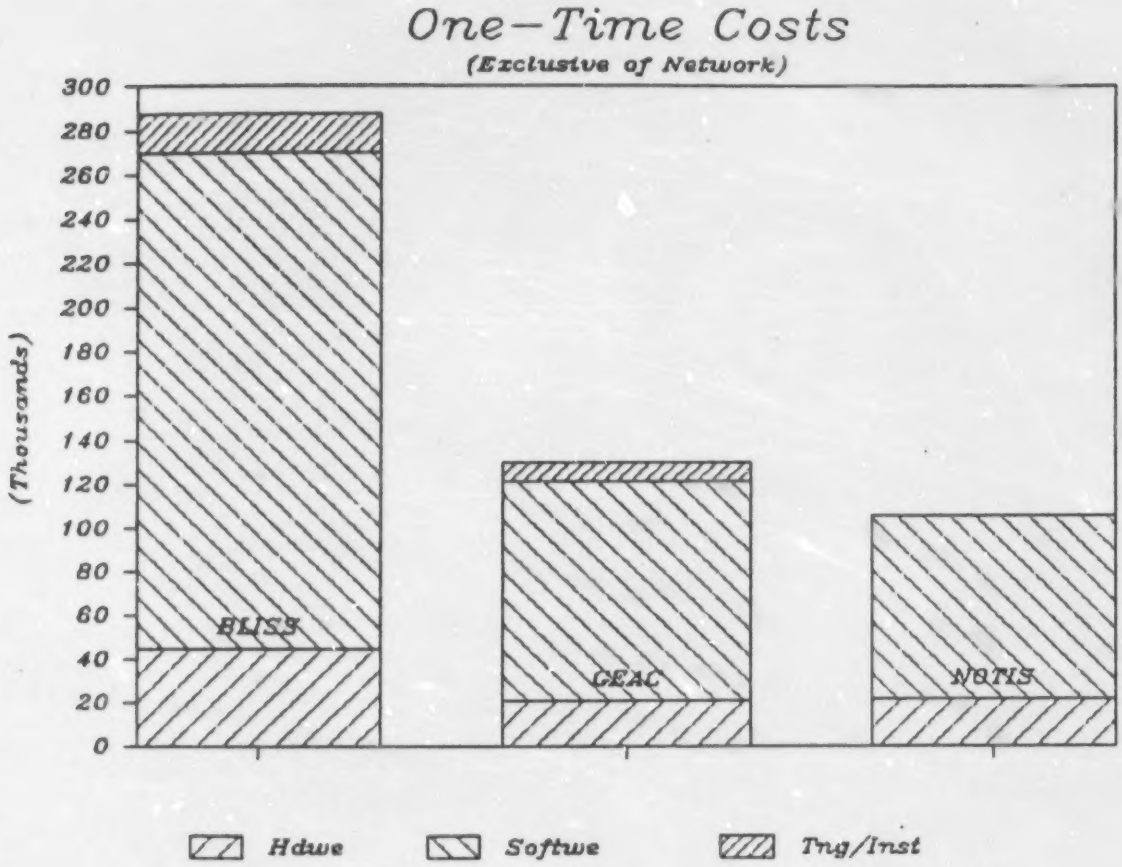
One-Time Costs
(Exclusive of Network)
(Thousands)
300
280
260
240
220
200
180
160
140
120
100
80
60
40
20
0
BLISS
CEAC
NOTIS
Hdwe
Softwe
Trng/Inst

Figure 15.

Capitalization of equipment is the clear choice over treating it as expense, since it appears that periodic upgrading to increase capacity will be necessary. The expense line is held constant except for an inflation factor, with price-performence improvements providing the incremental capacity. [2] Figure 16 compares recurring costs for the three recommended systems for an average library.

3.2 Total Institutional Costs

The last scenario presented in the preceding chapter produced annual income for the DRN as shown in Figure 14. From this, an annual percentage increase (or decrease) in DRN income can be computed. If a mean estimated annual RLIN billing is calculated for fiscal 1984 using five "average" RLG member institutions, then the behavior of this number in succeeding years can be approximated using the annual percentage change in DRN income.

Total annual recurring costs for the local system are estimated by taking the costs for a GEAC system (as an example) and adding the estimates for the network components using the heuristics mentioned above. The installation of the local system is assumed to occur at the beginning of fiscal 1987.

A component allowing for direct terminal communication with the DRN is also present: twelve terminals are assumed at the outset, going to three with the advent of the local system. The remaining terminals will be used in conjunction with those applications and enhancements not present on the local system (e.g. Archives and Manuscripts, Special Data Bases). The cost of the nine de-installed terminals is partially displaced by a dedicated circuit for machine/machine communications. All DRN communication costs are inflated using the Stanford Long Range Forecast factor.

The Program and Partnership Fee is likewise included. It is also inflated using the Stanford Long Range Forecast factor.

3.3 Conclusions

Figure 17 portrays the behavior of institutional library automation costs in a hypothetical RLG member institution based upon the average statistics. In this example, costs increase by a factor of 2.5 over eight years. In all likelihood, this factor is a lower bound, since in some cases increases in local system use will outpace the configuration growth pattern assumed.

---

2. See Carnegie Study on Distributed Computing, Part V, pp. 15-30

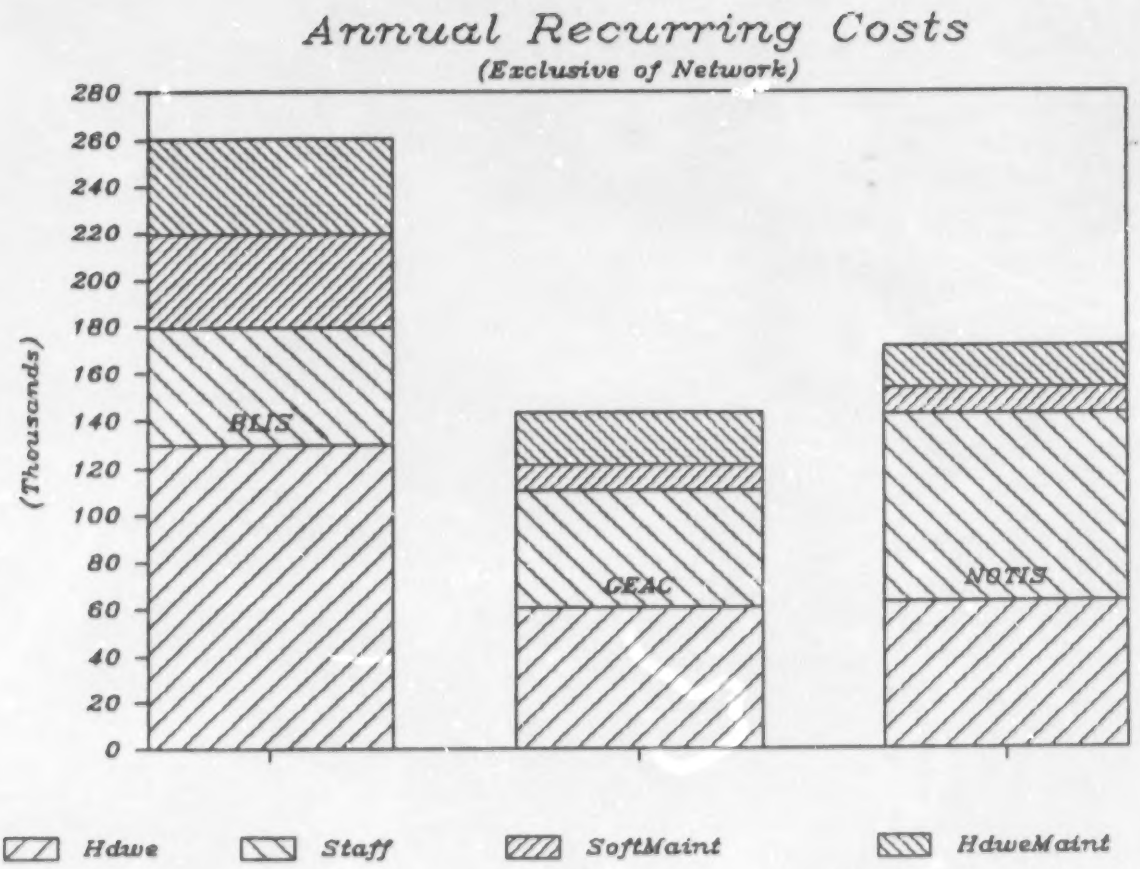
Annual Recurring Costs
(Exclusive of Network)
(Thousands)
280
260
240
220
200
180
160
140
120
100
80
60
40
20
0
BLIS
CEAC
NOTIS
Hdwe
Staff
SoftMaint
HdweMaint

Figure 16.

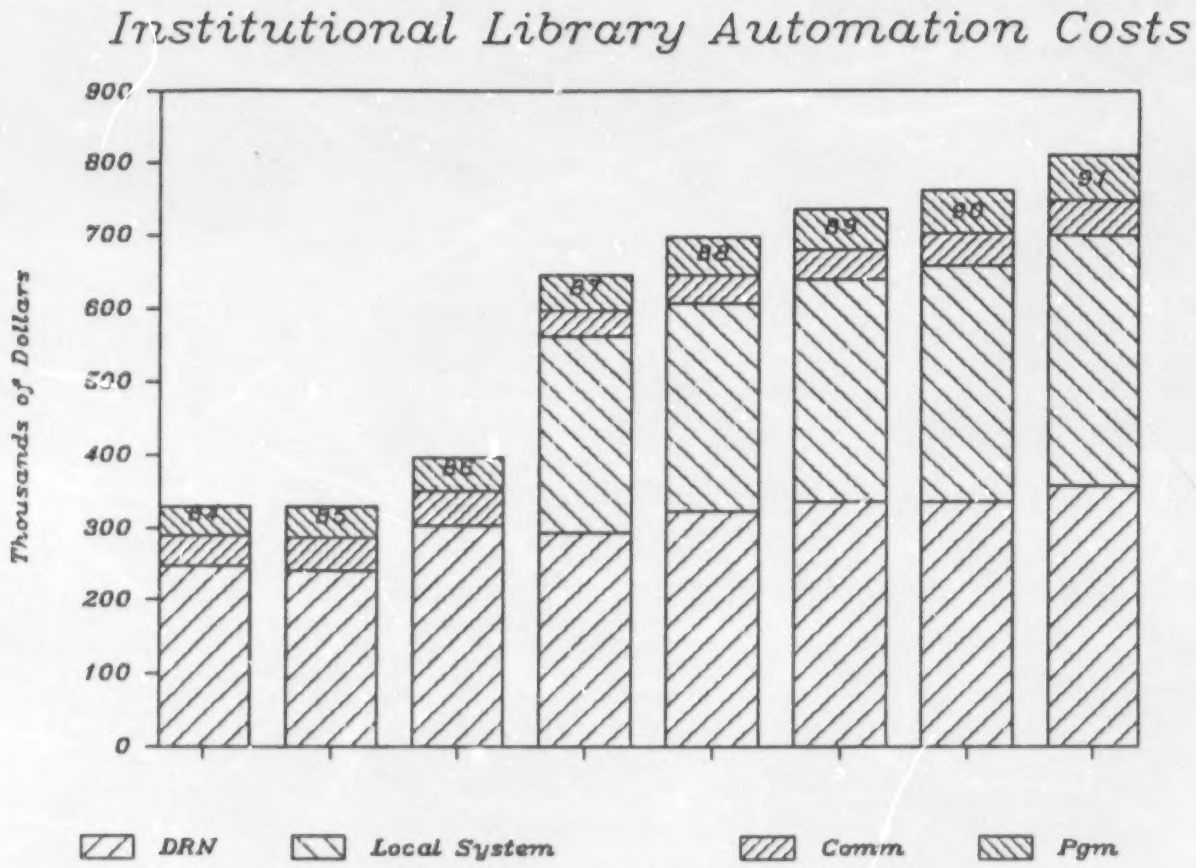
Institutional Library Automation Costs
Thousands of Dollars
900
800
700
600
500
400
300
200
100
0
84
85
86
87
88
89
90
91
DRN
Local System
Comm
Pgm

Figure 17.

It is emphasized that these costs are institutional. We must remind ourselves that use of the DRN in future years may well come from sources outside the library -- from the users themselves. The decision whether or not to have the library mediate access or not is a policy decision, not a technical one. If the decision to encourage unmediated access is taken, then costs will doubtless be offset by funds allocated to end users in their departmental budgets, from grant budgets, as well as from the automation line in the library budget. It is clear that institutional discussions of funding sources, fund allocation, and charge-back mechanisms should begin now. The library should serve as the catalyst for these discussions.

Appendix A

Detailed Scenario Assumptions

DISTRIBUTED PROCESSING
FINANCIAL MODEL ASSUMPTIONS

EXPENSE

Salaries & Benefits

* Inertial - Current Base (FY84 Budget of 43.56 FTE) holding constant; current benefit rate of 22.8% and vac. accrual 10.5% as stated in SLRF; dollars inflated at SLRF rate of 6.5% FY85 - FY91.

* Conservative - Staffing is reduced in the Lib Operations div as libraries leave the central system. All else remains according to inertial assumptions.
  FY85 -1 FTE
  FY86 -2 FTE
  FY87 -6 FTE
  FY88 -1 FTE

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions.

Professional Services

* Inertial - Increased by SLRF (5% FY85-FY91)

* Conservative - Reduction of $32k for Spires in FY86 out.

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions.

Computer Services

* Inertial -
  - Equipment -3081 - status quo - no buy out, payments continue assumes 0 FY88 through FY91
  - Disks - pay current IPA and IPA all incremental at current prices.
  - Drums - continue to lease upgrades and increments as required.
  - Printing - current lease plan through FY91
  - Shared Eqpt - FY84 cost thru FY91
  - K-Upgrade - 4 yr purchase plan beginning FY84
  - Maintenance - increased at SLRF - 7%
  - Licenses - increased at SLRF - 7%
  - Purch Svcs - increased at SLRF - 7%
  - CIT-Alloc Dir - increased at SLRF - 7%
  - CIT-Admin - increased at SLRF - 7%
  - CIT-Indir - 29% of the sum of CIT Alloc Dir + CIT Admin + Shared Eqpt + Production Control

* Conservative -
  - Equipment -3081 - Move to a 14 MIP uniprocessor at the beginning

of FY85, and run it through FY91. Four year purchase plan.

- Disks - pay current IPA and lease all incremental
- Drums - continue to lease upgrades and increments as required.
- Printing - current lease plan through FY91
- Shared Eqpt - FY87 goes to 0 through FY91
- K-Upgrade - IPA paid off in FY84, FY85 - FY91 goes to 0
- Maintenance - increased at SLRF - 7%
- Licenses - increased at SLRF - 7%
- Purch Svcs - increased at SLRF - 7% through FY91
- CIT-Alloc Dir - reduced $300k in FY87 out
- CIT-Admin - increased at SLRF 7%
- CIT-Indir - 29% of the sum of CIT Alloc Dir + CIT Admin + Shared Eqpt + Production Control

* Conservative Prime - Conservative assumptions

* Aggressive -
  - Equipment -3081 - purchase 14 MIP in FY85 and increase to 24 MIP at a constant K expense level thru FY91
  - Disks - pay current IPA and acquire increments via IPA until FY87 then acquire all new technology at lower prices.
  - Drums - lease thru FY86 then replace entire complement of new technology at reduced price/performance.
  - Printing - current lease plan through FY91
  - Shared Eqpt - FY87 cost goes to 0 through FY91
  - K-Upgrade - buyout in FY84; 0 thereafter
  - Maintenance - increased at SLRF - 7%
  - Licenses - increased at SLRF - 7%
  - Purch Svcs - prod ctrl assumed 0 in FY87 out; all other categories increased at SLRF 7%
  - CIT-Alloc Dir - reduced $300k in FY87 out
  - CIT-Admin - increased at SLRF 7%
  - CIT-Indir - 29% of the sum of CIT Alloc Dir + CIT Admin + Shared Eqpt + Production Control

Utilities
---

* Inertial - assumes FY85 base increase of $20k from FY84 increase of SLRF (8.2%) thereafter

* Conservative - inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - inertial assumptions

Production Supplies
---

* Inertial - assumes FY84 cost per activity unit (cat + acq) and assumes a FY84 base for 9700 supplies; rates increased by SLRF 5%

* Conservative - inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - inertial assumptions

Travel
---

* Inertial - increased at SLRF - 7%, assumes current staffing

* Conservative - reduced according to staff reductions; calculated by #libs at $/trip.

* Conservative Prime - Conservative assumptions

* Aggressive - inertial assumptions

Recruitment
---

* Inertial - increased at SLRF - 5%

* Conservative - increased at SLRF to FY87; 0 FY87 - FY91

* Conservative Prime - Conservative assumptions

* Aggressive - inertial assumptions

General Supplies
---

* Inertial - increased at SLRF - 5%

* Conservative - increased at SLRF to FY87; FY87 - FY91 reduction of 50%

* Conservative Prime - Conservative assumptions

* Aggressive - inertial assumptions

Other
---

* Inertial - increased at SLRF - 5%

* Conservative - increased at SLRF - 5%

* Conservative Prime - Conservative assumptions

* Aggressive - increased at SLRF - 5%

Space
---

* Inertial - increased at SLRF - 5%

* Conservative - increased at SLRF - 5%

* Conservative Prime - Conservative assumptions

* Aggressive - increased at SLRF - 5%

Other Equipment
---

* Inertial - increased at SLRF - 5%
* Conservative - increased at SLRF - 5%
* Conservative Prime - Conservative assumptions
* Aggressive - increased at SLRF - 5%

Host Charges
---

* Inertial - SLRF rate of 11.12% of sum of applicable costs
* Conservative - inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - inertial assumptions

Interest
---

* Inertial - status quo - Ford, Carnegie, Stanford Loans plus member repayment
* Conservative - inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - inertial assumptions

Principal
---

* Inertial - status quo - principal pmts to Ford, Carnegie, stanford, member repayments
* Conservative - inertial assumptions
* Aggressive - inertial assumptions
* Conservative Prime - Conservative assumptions

INCOME

Inertial, Conservative, Aggresive Scenarios

Activity	Cataloging	Acquisitions	Recon
FY84	1480905	960642	187610
FY85	1215961	744745	1341640
FY86	1230311	589595	1390035
FY87	1052261	557525	1002700

FY88	1052261	557525	905380
FY89	1052261	557525	358880
FY90	1052261	557525	457880
FY91	1052261	557525	369980

Conservative Prime Scenario

Cataloging - FY85-FY91 add 200k to the above
Acquisitions - FY86 = 744745
Recon - FY85 = 292000, FY86 = 1341640, FY87 = 1390035, FY88 = 1002700
FY89 = 905380, FY90 = 358880, FY91 = 457880

Cataloging

* Inertial - Activity levels with cost increased at SLRF 5% same mix of activity is assumed as FY84, ITPS pricing is used in same ratio as FY84.

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Retrospective Conversion

* Inertial - Activity levels, off peak pricing, cost per unit increased at SLRF 5%

* Conservative - Inertial assumptions
* Conservative Prime - FY85 assumes FY84 projected activity plus subject recon

* Aggressive - Activity levels increased by 10% each year, off peak pricing, cost per unit increased at SLRF

Catalog Cards

* Inertial - Catalog activity, cards assumed at 12.45 per title rates increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Tapes

* Inertial - ratio of tapes per title equals .9 (cat+acq+recon) cost increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Acquisitions

* Inertial - Activity levels, cost per unit increased at SLRF of 5%, same mix of activity assumed as FY84
* Conservative - Inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - Inertial assumptions

Acquisition Forms

* Inertial - Activity levels, cost per unit increased at SLRF
* Conservative - Inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - Inertial assumptions

Cat/Acq Maintenance

* Inertial - number of maintenance units is assumed 60% (FY84) of total Cataloging and 86% acquisitions, method of pricing is that of FY84 (charging for above median transactions), cost per unit increased at SLRF 5%
* Conservative - Inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - Inertial assumptions

Technical Processing Resource Usage

* Inertial - Activity levels, method and ratios assumed as in FY84 with cost increased at SLRF 5%
* Conservative - Inertial assumptions
* Conservative Prime - Conservative assumptions
* Aggressive - Inertial assumptions

Discounts/Commissions

* Inertial - CLASS remains at FY84 level increased at SLRF, discounts are as in FY84-75k% of operating income taken at 8% per annum on prepayment balances
* Conservative - Inertial assumptions
* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Reference Resource

* Inertial - assumed ratio of Ref to TP Res. as in JPM Scenario cost increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions with 10% increase in activity

Interlibrary Loan

* Inertial - 10% increase in activity per year is assumed, cost increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Searching

* Inertial - FY84 increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Miscellaneous

* Inertial - FY84 increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

RRS/FRNS/Worksheets

* Inertial - FY84 increased at SLRF 5%

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Authorities Control
---

* Inertial - new product assumed in FY85 at $.13/CPU Sec increased at SLRF 5%
  FY84 0 CPU Hrs
  FY85 65CPU Hrs
  FY86 95.68CPU Hrs
  FY87-FY91 101.12CPU Hrs each year

* Conservative - Inertial assumptions

* Conservative Prime - Conservative assumptions

* Aggressive - Inertial assumptions

Remote TP Searching
---

* Inertial - new product assumed in FY85 at $.13 increased at SLRF 5%; assumes central system activity of FY84 and as libs leave system activity goes to Remote TP Searching

* Conservative - rate increased to $1.25 and increased at SLRF 5%
* Conservative Prime - assumes $0 in FY85

* Aggressive - Conservative assumptions

Referred Public Service Searching
---

* Inertial - new product assumed in FY86 at $.13 increased at SLRF 5%;
  FY86 246CPU hrs/yr
  FY87 369 " "
  FY88 369 " "
  FY89-FY91 492CPU hrs/yr

* Conservative - new product assumed in FY86 at $.13 increased at SLRF 5%;
  FY86 738CPU hrs/yr
  FY87-FY91 984 CPU hrs/yr

* Conservative Prime - FY86 150 hrs, FY87 400 hrs,
  CPU Resource rate is doubled

* Aggressive - new product assumed in FY86 at $13. increased at SLRF 5%;
  FY86 738 CPU hrs/yr
  FY87 1470 " "
  FY88 1700 " "
  FY89 2100 " "
  FY90 2400 " "
  FY91 2400 " "